Journal of Economic Perspectives—Volume 4, Number 3—Summer 1990—Pages 149–170

In Honor of David Kreps, Winner of the John Bates Clark Medal

Jean Tirole

David Kreps is the winner of the 1989 John Bates Clark Medal, given to "that American economist under forty who is adjudged to have made a significant contribution to economic thought and knowledge." He obtained a doctorate degree in Operations Research in 1975 under the supervision of Evan Porteus at Stanford University. He has taught at Stanford University's Graduate School of Business since then. I will try to explain where Kreps's contributions fit in the food chain of economic ideas, why they have a high nutritional content, and why they should be in the diet of the common economist.[1]

I will arbitrarily divide this review into three major themes of Kreps's work: individual decision making, game theory (foundational and applied) and asset pricing. My strategy is to focus on a small number of papers to highlight better the lines of research; thus, this review is far from being exhaustive. In keeping with the spirit of this journal, I will focus on the economic contributions of the papers; some readers will thus feel that I do little justice to the technical mastery that characterizes his work.

[1]In doing so, I will (unsuccessfully) try to mimic Kreps's model review [32] of the contributions of the 1987 John Bates Clark Medal winner, Sandy Grossman, which appeared in the Spring 1988 issue of this journal.

■ *Jean Tirole is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts.*

Individual Choice under Uncertainty

Kreps's early work focused on individual decision-making under uncertainty [2, 3, 4, 7, 8, 9, 10, 11, 34]. Of particular interest to economists are his contributions on the issue of when uncertainty is resolved. His work with Evan Porteus [7, 8] studies a class of preferences that allow decision makers to care not only about their consumption stream, but also about the time at which uncertainty about future consumption is resolved.

There are two reasons why a single decision maker may prefer to obtain information early; that is, may be "resolution-seeking." The first is trivial and familiar. Getting information early allows better decision making; for instance, the agent may save more today if income tomorrow appears likely to be low.[2]

The second reason for being resolution-seeking is "psychological," which is to say that a decision-maker may simply prefer to have uncertainty resolved more quickly, even if knowing this information offers no way of changing the future outcome. Conversely, it is also possible for a decision maker to prefer to wait and find out the outcome only when it actually happens. A student who tries to learn his or her grade hours or days before it is publicly announced is "resolution seeking." An individual who prefers not to know about terminal illness is "resolution averse" (all the more that such knowledge would induce the individual to shift consumption forward and spend more time with family and friends).

Economists are used to working with the class of von Neumann-Morgenstern preferences, which leaves out the psychological factor and relies only on an expected utility calculation. To paraphrase [7, p. 185], consider a situation in which the flip of a fair coin determines the agent's consumption flow over two periods. If the coin comes up heads, consumption in the two periods will be (4, 12); if it comes up tails, consumption will be (4, 0). In the standard economist's model, the decision maker does not care whether the coin is flipped before the first time period or before the second; in either case, consumption will be 4 in the first period with a 50–50 chance of 12 or zero in the second period. For instance, an individual with von Neumann-Morgenstern utility function $\sqrt{c_1 + c_2}$ (where c_t represents consumption in period t) has expected utility $(1/2\sqrt{16} + 1/2\sqrt{4}) = 3$ in either case. But flipping the coin at the beginning of the first period resolves the uncertainty earlier.

Kreps and Porteus develop a broader class of preferences for which the individual can be resolution-seeking or resolution-averse. An example of preferences in the Kreps-Porteus class is $E_1(E_2(\sqrt{c_1 + c_2}))^{\alpha}$. E_t represents the

[2]But beware: this reason for preferring early resolution holds only when the agent is playing against nature. In a situation with multiple decision makers, receiving information may hurt the agent because other decision makers are aware of it and may strategically modify their behavior.

David Kreps

expectation operator for expectations formed at the start of periods 1 and 2, and α is a positive number which measures the degree of resolution-seeking behavior. When $\alpha = 1$, preferences fit the von Neumann-Morgenstern expected utility concept.

Now imagine that $\alpha = 2$, and the coin is tossed at the start of period 1. Then expected utility will be: $1/2\ (\sqrt{16})^2 + 1/2(\sqrt{4})^2 = 10$. By contrast, imagine that the coin is tossed at the start of period 2. In this case, expected utility will be $(1/2\sqrt{16} + 1/2\sqrt{4})^2 = 9$. Notice that in the first case, the two possible levels of total consumption are each squared before being given their 1/2 weights; in the second case, the two possible levels are combined with their 1/2 weights before being squared. In this way, immediate resolution is given more weight than postponed resolution. The reader should be able to affirm that individuals are resolution-seeking when $\alpha > 1$ and resolution-averse when $\alpha < 1$.

As we will see, the thread that unifies most of Kreps' work is that a dynamic context differs from and is richer than a static context. A von Neumann-Morgenstern decision maker analyzes complex dynamic asset trading strategies in terms of their random streams of associated consumptions and compares their expected utilities today; in that sense, classical choice theory reduces a dynamic problem into a static one. In contrast, Kreps and Porteus apply the expected utility approach only to immediately resolving lotteries whose prizes are composed of later-resolving lotteries, and connect preferences at various dates with a temporal consistency axiom that guarantees that decision makers do not take decisions today they will regret tomorrow. As in the example above, preferences are then represented by a sequence of von Neumann-Morgenstern utilities rather than a single one.

Clearly, preferences are no longer defined by the stream of consumptions. A resolution-seeking decision maker might prefer lower levels of consumption, with early knowledge of the exact levels, to higher levels of consumption where knowledge about the exact level is delayed.

Researchers have recently noted that such preferences might be used to explain empirical anomalies in the elasticities of intertemporal substitution and in the equity premium. Many studies have estimated models where a representative consumer has time-separable von Neumann-Morgenstern preferences, and they have generally estimated that intertemporal elasticities of substitution are very small. This finding is troublesome, because the coefficient of relative risk aversion (measuring aversion to the income risk within a period) is necessarily equal to the inverse of the intertemporal elasticity of substitution. Thus, very small intertemporal elasticities of substitution mean an unreasonably high coefficient of relative risk aversion. Similarly, it seems difficult to find plausible discount factors and coefficients of relative risk aversion that explain the levels of real rates of interest and the large mean premia on equities over

Table 1
Works by David M. Kreps

[1] "A Note on Fulfilled Expectations' Equilibria," *Journal of Economic Theory*, February 1977, *14*, 32–43.
[2] "Decision Problems with Expected Utility Criteria. I: Upper and Lower Convergent Utility," *Mathematics of Operations Research*, 1977, *2*, 45–53.
[3] "Decision Problems with Expected Utility Criteria. II: Stationarity," *Mathematics of Operations Research*, 1977, *2*, 166–174.
[4] "On the Optimality of Structured Policies in Countable Stage Decision Processes. II: Positive and Negative Problems," (with Evan L. Porteus), *SIAM Journal on Applied Mathematics*, 1978, *16*, 420–428.
[5] "Sequential Decision Problems with Expected Utility Criteria. III: Upper and Lower Transience," *SIAM Journal on Control and Optimization*, 1978, *16*, 420–428.
[6] "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," (with J. Michael Harrison), *Quarterly Journal of Economics*, May 1978, *92*, 323–336.
[7] "Temporal Resolution of Uncertainty and Dynamic Choice Theory," (with Evan L. Porteus), *Econometrica*, January 1978, *146*, 185–200.
[8] "Temporal von Neumann-Morgenstern and Induced Preferences." (with Evan L. Porteus), *Journal of Economic Theory*, February 1979, *20*, 81–109.
[9] "Dynamic Choice Theory and Dynamic Programming," (with Evan L. Porteus), *Econometrica*, January 1979, *47*, 91–100.
[10] "A Representation Theorem for Preference for Flexibility," *Econometrica*, May 1979, *147*, 565–578.
[11] "On Sophisticated Choice of Opportunity Sets," Cambridge University Economic Theory Discussion Paper No. 8., 1979.
[12] "Martingales and Arbitrage in Multiperiod Securities Markets," (with J. Michael Harrison), *Journal of Economic Theory*, June 1979, *20*, 381–408.
[13] "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities," *Journal of Mathematical Economics*, March 1981, *8*, 15–35.

Table 1 *(continued)*

Works by David M. Kreps

[14] "Multiperiod Securities and the Efficient Allocation of Risk: A Comment on the Black-Scholes Option Pricing Model." In McCall, John J., ed., *The Economics of Information and Uncertainty*. Chicago: The University of Chicago Press, 1982.

[15] "Asset Pricing Models," draft, 1982.

[16] "Reputation and Imperfect Information," (with Robert Wilson), *Journal of Economic Theory*, August 1982, *27*, 253–279.

[17] "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," (with Paul Milgrom, John Roberts, and Robert Wilson), *Journal of Economic Theory*, August 1982, *127*, 245–252.

[18] "Sequential Equilibria," (with Robert Wilson), *Econometrica*, July 1982, *50*, 863–894.

[19] "Models in Managerial Accounting," (with Joel Demski), *Journal of Accounting Research*, Supplement 1982, *20*, 117–148.

[20] "Quantity Pre-Commitment and Betrand Competition Yield Cournot Outcomes," (with Jose Scheinkman), *Bell Journal of Economics*, Autumn 1983, *14*, 325–337.

[21] "Modeling the Role of History in Industrial Organization and Competition," (with A. Michael Spence). In Feiwel, G., ed., *Issues in Contemporary Micro-economics and Welfare*. London: Macmillan and Company, 1985, 340–378.

[22] "Models of Spillovers in R & D," (with A. Michael Spence). In Tsuchiya, M., ed., *Technological Innovation and Business Strategy*. Tokyo: Nippon Keizai Shimbunsa Press, 1986 (in Japanese; English version available in draft form).

[23] "Corporate Culture and Economic Theory," In Tsuchiya, M., ed., *Technological Innovation and Business Strategy*. Tokyo: Nippon Keizai Shimbunsa Press, 1986 (in Japanese; English version reprinted in Alte, J., and K. Shepsle, eds., *Rational Perspectives on Positive Political Economy*, Chapter 4, Cambridge: Cambridge University Press, 1990, pp. 90–143).

[24] "Price Destabilizing Speculation," (with Oliver D. Hart), *Journal of Political Economy*, 1986, *94*, 927–952.

[25] "Three Essays on Capital Markets," *La Revista Espanola de Economia*, 1987, *4*, 11–145.

[26] "Rational Learning and Rational Expectations," (with Margaret Bray). In Feiwel, G., ed., *Kenneth Arrow and the Ascent of Economic Theory*. London: Macmillan 1987.

[27] "Signaling Games and Stable Equilibria," (with In Koo Cho), *Quarterly Journal of Economics*, 1987, *102*, 179–221.

[28] "Consistency, Structural Consistency, and Sequential Rationality," (with Garey Ramey), *Econometrica*, 1987, *55*, 1331–1349.

[29] "Nash Equilibrium," *The New Palgrave*. London: Macmillan, 1987.

[30] "Reputation in the Simultaneous Play of Multiple Opponents," (with Drew Fudenberg), *Review of Economic Studies*, 1987, *54*, 541–568.

[31] "On the Robustness of Equilibrium Refinements," (with Drew Fudenberg and David Levine), *Journal of Economic Theory*, 1988, *44*, 354–380.

[32] "In Honor of Sandy Grossman, Winner of the John Bates Clark Medal," *Journal of Economic Perspectives*, Spring 1988, *2*, 111–135.

[33] "Out of Equilibrium Beliefs and Out of Equilibrium Behavior," draft, 1986. In Hahn, F. H., ed., *The Economics of Information, Missing Markets, and Games*. Oxford: Oxford University Press, 1989.

[34] "Intertemporal Preferences with a Continuous Time Dimension: An Exploratory Study," (with Chi-fu Huang), draft, 1987.

[35] "The Folk Theorem with Long and Short Run Players," (with Drew Fudenberg and Eric Maskin), draft, 1987.

[36] "Static Choice in the Presence of Unforseen Contingencies," draft, 1988.

[37] "Signaling," (with Joel Sobel), draft, 1988. To appear in Aumann, R., and S. Hart, eds., *The Handbook of Game Theory*.

[38] *Notes on the Theory of Choice*. Boulder: Westview Press, 1988.

[39] *A Course in Microeconomic Theory*. Princeton: Princeton Universty Press, 1990.

[40] *A Theory of Learning, Experimentation and Equilibrium in Games*, (with Drew Fudenberg), partial draft, 1988.

the riskless rate.[3] Models of "habit formation" such as Constantinides (1988) attempt to resolve this problem while keeping both the representative consumer approach and von Neumann-Morgenstern preferences, by not allowing utility of consumption to be time-separable. In this case, determining the utility of consumption in one period requires knowing consumption in all other periods. Another route, taken by Epstein-Zin (1989), Farmer (1990), and Weil (1989, 1990), consists of positing Kreps-Porteus preferences for the representative consumer. The intertemporal elasticity of substitution and the coefficient of relative risk aversion are now disconnected.

Kreps [10] also extended the standard model of intertemporal decision making by formalizing the notion of preference for flexibility. A standard reason (Henry, 1974; Goldman, 1974) why decision makers prefer to have options (a large opportunity set) is that they may not know their preferences in advance, since they will learn about the world in the future. In the von Neumann-Morgenstern framework, this situation can be depicted by state-dependent preferences, where utility depends both on the amount of personal consumption and on the state of nature in which this consumption is realized. The decision maker can adjust consumption to fit his or her preferences by not committing to a specific consumption pattern in advance and waiting until the state of nature is learned. For instance, the decision maker may delay irreversible investment to have a larger opportunity set tomorrow.

Instead of positing states of nature as in the traditional approach and analyzing preferences over state-contingent outcomes, Kreps [10] defines preferences over opportunity sets, from which the decision maker will later choose a single object; for instance, the preferences are over menus, from which a meal will later be chosen, rather than over state-contingent meals. He assumes that (a) preferences are transitive and complete (a strong assumption: the decision maker is able to compare any two opportunity sets); (b) there is a desire for flexibility (a larger opportunity set is weakly preferred); and (c) if the decision maker is indifferent between two nested opportunity sets, enlarging each of these opportunity sets by adding the same feasible outcomes preserves indifference.[4] The main theorem is that having preferences over opportunity sets satisfying these assumptions (plus some other reasonable ones) is equivalent to behaving as if there were underlying states of nature, in spite of the fact that no explicit uncertainty exists. In [36], Kreps argues that this representation theorem may eventually shed light on some aspects of multi-person choice when

[3] Bewley (1982) and others would argue that the problem is that the representative consumer approach biases the estimates of the coefficient of risk aversion when consumers are heterogeneous and markets are incomplete.

[4] Indifference between two nested opportunity sets means that the decision maker attaches "no value" to the extra opportunities contained in the bigger set. To see why the sets must be nested in this assumption, consider for instance two disjoint sets between which the decision maker is indifferent. Adding the outcomes of the first set (or more generally substitute outcomes) does not change this set but in general makes the second set more attractive.

states of nature are not fully described, a notion prominent in the transaction cost analysis of Williamson (1975, 1985) and Grossman and Hart (1986).

Game Theory

David Kreps, like several prominent theorists of his generation including Bengt Holmström, Paul Milgrom, Bob Rosenthal, and Al Roth, crossed paths with Robert Wilson, and their collaboration spurred Kreps to a new and fruitful course of research. I will first discuss Kreps' work in defining new equilibrium concepts ([18], [27], [28]), and then his expression of doubts about some refinements ([28], [31], [33]) and his learning approach to equilibrium behavior ([26], [40]). His work in game theory is characterized by a creative tension between the design of new refinements of Nash equilibrium and a critical attitude toward them. I will then discuss his applications of game theory to particular classes of games.

Foundations of Game Theory

Before exploring Kreps and Wilson's [18] notion of sequential equilibrium, which is a core concept for analyzing dynamic games of incomplete or imperfect information, a cursory overview of the development of non-cooperative game theory is helpful.

A Nash (1950) equilibrium of a game is a set of strategies, one for each player, such that each player's strategy is an optimal response to the strategies of the other players; in equilibrium, no player has a strategy that yields a higher expected payoff. The first and most familiar example is Cournot's duopolists who simultaneously choose their output to maximize their individual profits. But a strategy can be more complex than the choice of a single action. It may more generally depict the player's contingent behavior at different dates for different histories of the game or information for the player. Harsanyi (1967–68) generalized the notion of Nash equilibrium to games of incomplete information, that is, to games in which players have private information about payoffs. He offered to view each player as a set of "incarnations" where each incarnation corresponds to a different private information. Incarnations of the same player in general behave differently because they receive different payoffs. For instance, in the Cournot game, a firm with high cost will choose a lower output than the same firm when its cost is low. In a Bayesian equilibrium, players do not know which incarnations of the other players they are playing against, and each tries to maximize his or her expected payoff given private information, and given private beliefs about the probability distribution of rivals' incarnations.

In a contemporary development, Selten (1965) pointed out that the Nash concept may not be restrictive enough in situations in which the players do not choose actions simultaneously and once and for all.[5] Suppose that firm 1 sets its output before firm 2, and that firm 1 chooses a quantity level which it knows to be a (Cournot) Nash equilibrium, were both firms to be moving simultaneously. Firm 2 chooses the strategy of setting its quantity equal to the Cournot equilibrium level if firm 1 does so, but threatens to choose a very large quantity if firm 1 deviates from the Cournot equilibrium level. Neither of the firms can raise its profits by changing its strategy. If firm 1 produces a different level, it suffers from a surge of production from firm 2. And, since firm 1 chooses the Cournot equilibrium output, so does firm 2.

Yet this Nash equilibrium seems suspicious: were firm 1 to deviate from its Cournot equilibrium output, firm 2 would be better off deciding on an optimal reaction, rather than flooding the market. The problem with this Nash equilibrium is that it allows firm 2 to make "empty threats." With the equilibrium strategies as specified, such threats are costless, since the party making the threat knows for sure that it will never have to carry it out. (In the parlance of game theory, firm 1's outputs that differ from the Cournot level are "off-the-equilibrium-path" for this Nash equilibrium. The goal of the "refinement-of-Nash-equilibrium" literature is to put restrictions on what players do if their opponents do something totally unexpected.) Selten's (1965) (subgame) perfect equilibrium concept requires that the players' strategies be optimal for any history of the game, and not only histories that arise on the equilibrium path. The idea of subgame perfection is just that the strategies specified in any contingency must themselves be credible, in the sense that a rational player would carry them out should the contingency ever arise.

In multi-period games in which players have private information (as in Harsanyi), a player's optimal strategy for some history of the game depends on his or her (posterior) beliefs about the opponents' private information. Posterior beliefs in turn should be updated from prior beliefs using the information conveyed by past observed behavior and presumed (equilibrium) strategies. Selten (1975) introduced the concept of trembling hand perfect equilibrium that implicitly combines subgame perfection and Bayesian equilibrium.[6] He refined the idea of Nash equilibrium by examining the case where players "tremble" when choosing their strategies; that is, they are most likely to choose their preferred strategy, but retain a small chance of playing any other strategy. The technical trick of introducing trembles is that every contingency will sometimes happen, and thus all threats are costly because there is a chance they will need to be carried out. For instance, in the sequential duopoly game, firm 2 plays its optimal reaction for any output of firm 1 because all outputs of firm

[5]This point was already made in a simple duopoly game by von Stackleberg, who changed the timing of the Cournot quantity setting game.

[6]The learned reader will note that the following description of Selten's work is not quite accurate. The "agent normal form," a central piece in Selten's construction, incorporates information about the extensive form by identifying "incarnations" of a single player who moves at different information sets (for example, different dates).

1 have at least a small probability of being played. A more subtle implication of trembles is that Nash behavior requires implicitly players to update their beliefs about each others' private information using Bayes rule. (Bayes rule always uniquely determines posterior beliefs as any history has positive probability.) A trembling-hand perfect equilibrium is a limit of Nash equilibria in which the players' trembles on suboptimal strategies vanish.

While Selten's motivation is dynamics, his construction focuses on strategies and updating of beliefs remains implicit. One contribution of Kreps and Wilson [18] was to stress the importance of beliefs, which become an explicit part of the description of equilibrium. In doing so, they make Selten's point that dynamic games are more than static games, with a vengeance. An "assessment" is a pair of strategies and beliefs at each "information set." In most economic games, this amounts to having strategies and beliefs be defined at each stage of the game. This framework makes beliefs an explicit part of the description of an equilibrium.

Kreps and Wilson first require that strategies be "sequentially rational," which, in a multi-period context, means that the strategies starting at any point in time form a Bayesian equilibrium given the beliefs at the beginning of the period.

Kreps and Wilson also require consistency of beliefs. An equilibrium assessment must be the limit of assessments in which all strategies are played with positive probability and Bayes rule is applied.[7] Thus, another difference with trembling hand perfection is that the strategies are required to be in equilibrium only in the limit and not necessarily along the sequence of trembles, which makes checking for equilibrium somewhat easier. While this implies that there are at least as many sequential equilibria as trembling hand perfect equilibria, Kreps and Wilson show that the two concepts yield the same set of equilibrium outcomes for almost all games. They also thoroughly analyze the structure of equilibria. For example, they show that for almost all games with finite numbers of players and actions, the set of sequential equilibrium outcomes is finite.

By making beliefs a central piece of the definition of equilibrium, Kreps and Wilson [18] sparked a sizeable literature on how to constrain beliefs about events that have zero chance of happening. Kreps, together with Cho [27] led the way by investigating some restrictions in signaling games. Signaling games

[7] Kreps and Wilson discuss the behavioral implications through a series of examples. For multi-stage games of incomplete information (in which players have private information about their preferences), sequential equilibrium is simply equivalent to "perfect Bayesian equilibrium," which requires (1) having a Bayesian equilibrium from each period on given the beliefs in that period, (2) updating beliefs from one period to the next using Bayes rule any time the actions have positive probability and (3) not letting players signal information that they don't possess (that is, the beliefs about player 1's information are not affected by player 2's actions, except when the two players' informations are correlated, in which case the change in beliefs about player 2's information induces an indirect change in beliefs about player 1's information). See Fudenberg and Tirole (forthcoming).

are two-player games in which one player has private information and moves first. The other player(s), who moves second, may or may not have private information and is affected by the first player's information.

The archetypical signaling game is Spence's (1974) education model in which a worker invests in education partly to signal an ability to two employers who then bid for the worker's services. Consider the following example: an employee has a 50–50 chance of being a low or high ability worker. Output q results directly from ability a: $q = a$, but there is no way to draw up an enforceable contract based on output. Note that productivity does not depend on education in this example. However, the worker can get some level of education to signal his or her ability level. Let us say that the worker's utility is equal to $w - e/a$, where w is the wage rate received, e is the level of education, and a is the ability level. Notice that the higher the ability level, the less a marginal investment in education will reduce utility for a given wage. Thus, the marginal cost of signaling through education decreases as ability rises.

If the competitive employers knew the worker's ability, the worker would receive a wage equal to ability level, and would not invest in education. If the employers have incomplete information about the worker's ability, many sequential equilibria exist depending on the way the educational signal is interpreted. For instance, consider the equilibrium called the "least cost separating equilibrium" or LCSE. Say that a low ability worker has an ability level of 1; chooses an education of zero; and receives wage level of 1. A high ability worker has an ability level of 2; chooses an education level of 1; and receives a wage level of 2. As long as the level of education is at least 1, then a low-ability worker has no incentive to try to signal that he or she has high ability (since $1 - 0 = 2 - 1$). This sequential equilibrium might be sustained by the employers' belief that a worker who chooses to invest something other than zero or 1 in education must be low ability.[8] It results in low- and high-ability workers having utilities of 1 and 1.5 respectively.

The socially efficient equilibrium, on the other hand, has both types choose an education level of zero and receive a wage of 1.5, representing the 50–50 probability that the worker has an ability level of 2 or of 1. To sustain this equilibrium, employers may have the belief that any worker who chooses to be educated is low ability (see previous endnote). There are many other sequential equilibria, corresponding to a range of beliefs about how the educational signal is to be interpreted.

[8] In the parlance of the previous discussion, education levels other than those chosen by some type of worker (that is, differing from 0 and 1 in the LSCE) are off-the-equilibrium path levels. Because they have probability zero, Bayes rule does not pin down the employers' posterior beliefs about the worker's ability following such levels. Note that the trembles associated with sequential equilibrium do not help in this respect. By appropriately choosing the ratio of the probabilities that the low- and high-ability types tremble on an off-the-equilibrium-path education level, the modeler can generate any posterior beliefs. See also note 7 on intuition why beliefs are fairly unconstrained by sequential equilibrium in such games.

Cho and Kreps view the LCSE as the "most reasonable" sequential equilibrium. In particular, they argue that the socially efficient equilibrium cannot sustain itself, based on the following reasoning. Recall that both sides choose zero education and have a utility of 1.5. A high ability worker could choose education $e = .6$ and give the following "speech" to the employers (although speeches are not explicitly part of the game): "If I were the low-ability type, I would not choose an education level of .6, because I would get at best $2 - .6/1 = 1.4$, which is lower than my equilibrium utility, while if I am the high ability type, and you believe this speech, I will gain from deviating as $2 - .6/2 = 1.7 > 1.5$." Thus, the employers should infer that a worker who chooses an education level of .6 is high ability and one who does not is low ability, which destroys the pooling equilibrium. Stiglitz argued that if this speech is indeed credible, and the level of education .6 separates those of high and low ability, then low-ability types who are now receiving wages equal to 1 will have an incentive to choose an education level of .6 after all. Cho and Kreps [27] discuss this issue further. Cho and Sobel (forthcoming) and Banks and Sobel (1987) provide further elaborations on the Cho-Kreps refinement. Kohlberg and Mertens (1986) offer a different philosophy about refinements.

The equilibrium selected by the Cho-Kreps criterion in this education game exhibits an interesting discontinuity in the effect of beliefs: as long as both low- and high-ability types exist, the same least cost separating equilibrium holds. That is, the probability of a low-ability worker could be any small fraction, rather than 1/2, without any effect on equilibrium behavior. As the probability of a low-ability worker shrinks toward zero, a high-ability worker still invests a substantial amount in education ($e = 1$) even though the employers are very confident that the worker is of high ability. However, when it is certain that the worker is high ability, the worker does not need to invest in education to obtain a high wage level ($w = 2$). This example suggests that adding new types can introduce very different equilibria. In [31], Fudenberg, Kreps and Levine examine this issue in general games.[9]

With Fudenberg, Kreps has launched an ambitious and interesting project on the learning foundations of various equilibrium concepts. One interpretation of equilibrium behavior is that players know the structure of the game and

[9]They show that any pure strategy Nash equilibrium of a game is close to a "strict" equilibrium of a perturbed game in which one has introduced new types with small probabilities. (Mixed strategy equilibria must obviously be purified to apply the Fudenberg-Kreps-Levine methodology.) A strict equilibrium is an equilibrium in which each player has a unique best response to rivals' strategies (a strong concept in static games, ruling out mixed strategies, and even more in dynamic games, ruling out zero-probability histories); but it suffices for the reader to know that a strict equilibrium satisfies the Cho-Kreps or even the Kohlberg-Mertens refinements. Among other things, this result can be interpreted as a criticism of the class of equilibrium refinements studied in the late 1980s; if one believes in such refinements, it is hard to discard Nash equilibrium strategies that they do not select, as they (or rather nearby strategies) do satisfy these refinements in a slightly perturbed version of the game and the modeler cannot be quite certain about the exact information structure of the game.

introspectively compute equilibrium. Another potential justification of equilibrium is that players learn about their opponents' behavior from observing their opponents past behavior. In [40], players are boundedly rational.[10] First, the players use adaptive expectations to predict their rival's behavior. Second, the players sometimes experiment. They maximize their current expected payoff with high probability, but not with certainty. The rate of experimentation decreases over time but not so quickly as to generate insufficient information about out-of-equilibrium play (experiments on any action at any information set must occur infinitely often).

Whether Nash equilibria (or refinements thereof) develop in the long run depends on the way players experiment. Like the trembles of sequential equilibrium, experimentation supplies a complete interpretation of out-of-equilibrium behavior. It is shown that strategies which are not Nash equilibria are not locally stable, which means that there is no initial history of play close to the strategies such that the sequence of learning yields these strategies in the long run with probability close to one. Conversely, if an outcome is a Nash equilibrium outcome, then it is locally stable for some learning behavior satisfying their assumptions (for every learning behavior if the Nash equilibrium is "strict," that is, is such that each player has a unique optimal reaction to opponents' strategies).

To rule out Nash equilibria that do not satisfy a given refinement as limits of the learning process, one must put more structure on beliefs and on experimentation. For instance, to obtain something resembling sequential equilibrium, which requires that two players' beliefs about a third player's information or action must coincide (as they result from the same trembles), players must asymptotically agree not only on the probability of equilibrium actions, but also on the "relative probability" of asymptotically zero-probability actions, which requires further conditions. To obtain something resembling the Cho-Kreps criterion, one must put even more structure on experimentation. In the education example, in particular, the able worker must experiment much more often on high education levels than the less able worker.

Learning is actually an old preoccupation for Kreps. In [26], he and Bray take on rational rather than adaptive learning. Here agents use Bayesian updating to learn about their environment. [26] makes the simple, yet important observation that beliefs, being bounded martingales,[11] must converge. This convergence of beliefs is general and can be applied to games or to rational expectations equilibria.

[10]The precise assumptions are complex. We only give some of their flavor here.

[11]The expectation of tomorrow's posterior probability of a state of nature is equal to today's probability of this state of nature. Probabilities are bounded because they are between zero and one. The martingale convergence theorem implies that the (updated) probability of a state of nature, given a sequence of increasingly finer (or at least not coarser) information structures, converges to some number in the long run with probability one.

Kreps and Bray give an example in [26] of a repeated market for an asset with a one-period life. An agent receives imperfect information about the value of the current asset each period and trades with an uninformed agent. The uninformed agent also does not know the (time-invariant) coefficient of risk aversion of the informed agent (but has prior beliefs about it) and may learn it over time. One can imagine that each of the two agents represents a continuum of traders with the same preferences and information. Under some conditions, in an intertemporal rational expectations equilibrium, the uninformed agent eventually learns the risk aversion parameter and the economy converges to a static rational expectations equilibrium of the type described in Grossman and Stiglitz (1980). This latter property is very hard to generalize even if the environment is stationary and an intertemporal rational expectations equilibrium exists. Bray and Kreps offer a number of interesting reasons why the spot market equilibria may never settle down to some stationary rational expectations equilibrium (even though the uninformed agents' beliefs always converge).

Applied Game Theory

Some of Kreps's most widely read contributions examine the role of reputation. In [16], [17] and Milgrom-Roberts (1982a), what came to be known as the gang of four (Kreps-Milgrom-Roberts-Wilson) proposed using Bayesian games to depict the notion that players may want to develop a reputation for being fair, cooperative, tough or having whatever characteristic will induce friendly behavior by other players in the future. They offer two examples: the case of a chain-store that may want to prey on earlier entrants to deter further entrants [16] (extended in [30] by Fudenberg and Kreps to allow one of the parties to sustain a reputation against several opponents played simultaneously) and players in a repeated prisoner's dilemma who want to develop a reputation for being cooperative [17].

In a sense, most dynamic Bayesian games can be thought of as reputation problems. The example of a low-ability or high-ability worker is an instance of a player trying to develop or maintain a reputation for being able. Similarly, Milgrom and Roberts (1982b) describe an incumbent firm which charges a low price to convey the information (develop a reputation) that it has low cost and thus deter entry. Reputation models should be distinguished from general Bayesian models. Reputation models allow broad uncertainty about preferences. Both [16] and [17] allow players to be "crazy." The central problem in reputation models is to convince the opponents that you will do something irrational; for example, employ predatory pricing no matter what [16] or cooperate in an irrational manner [17]. Success in doing this will result in higher equilibrium payoff. The gang of four allows my opponents to think there is some chance that I am crazy. However, they assume that crazy types have small probability and emphasize long, repeated relationships, as opposed to most Bayesian models where the probabilities of types are generally larger and an informed party must sustain a reputation over a small horizon (remember, the high-ability, low-ability signaling game had only a one-period horizon). One of the interesting points in the analysis is that the number of repetitions of the game needed to obtain reputation phenomena (rational types pretending they are crazy) may be fairly small. For instance, in the predation game of [16], the minimum probability that the incumbent firm simply enjoys predatory behavior needed to sustain reputation decreases exponentially with the number of periods.

The gang of four's contributions set up a flurry of exciting applications to industrial organization, organization theory, macroeconomics, international economics, finance and other fields. In fact, Fudenberg and Maskin (1986) warn against an abusive use of models with large horizons and small probabilities of craziness. They show that "anything" (in technical terms, any individually rational payoff) can be sustained in a long (yet finite) repeated game when the discount factor is close to 1 and the modeler chooses what form of craziness will occur with at least a small probability.[12]

Thus, to salvage the reputation story from the Fudenberg-Maskin charge of arbitrariness, it is important to limit what sorts of craziness should be allowed. First, one can argue (with the gang of four) that information exists about the most likely forms of craziness: a predatory type in [16], a tit-for-tat person in [17], and so on. Second, and a departure from the spirit of reputation models, one can restrict attention to situations which can be formulated in terms of large uncertainty, rather than craziness. In this line of reasoning, the predatory type in [16] might actually be a rational firm that just happens to have very low costs and thus wants to charge low prices regardless of the presence of the entrant. Third, Fudenberg and Levine (1989) eliminate the indeterminacy of the Fudenberg-Maskin result in the case in which a long-term player faces a sequence of short-term players, by showing that patient long-term players are able to sustain whatever reputation they prefer,[13] when the horizon is sufficiently long. In this case, the choice of the form of craziness is made by a player in the game, not by the modeler. This result however does not help in removing indeterminacy in games with several long-term players, such as the repeated prisoner's dilemma.

[12]It is easy to prove a weaker form of their result, namely that any payoff vector above a Nash equilibrium payoff vector is sustainable when one introduces a small probability of craziness. Consider a two-player game, Nash equilibrium strategies s^* of the (one-period) stage game in the absence of crazy types, and strategies s of the stage game that yield the two players more than their Nash payoff. Suppose now that, with probability ε, each player is crazy in the following way: each player plays s_i as long as s has been played before, and s_i^* until the end of the horizon if anybody has deviated in the past. Players do not want to deviate from s if the discount factor is close to 1 and the number of remaining periods is sufficiently large because a deviation would mean reversion to Nash behavior and has cost approximately proportional to ε times the length of the remaining horizon, which exceeds the one-period gain from deviating.

[13]That is, the patient long-term player obtains approximately the payoff that would be obtained if her opponents were convinced that the player had the (possibly crazy) type that elicits from them the most favorable responses, assuming that this type has positive prior probability.

Let me conclude this section on applied game theory with two well-known articles by Kreps in the area of industrial organization. Written with Scheinkman, [20] goes back to the foundations of Cournot equilibrium. Cournot's description of oligopoly behavior is often (rightly) criticized for assuming the exstence of a price-setting auctioneer who clears the market given the quantities set by the firms. In practice, prices are chosen by the firms and the best interpretation of the Cournot model is that the quantities stand for the production or sale capacities previously set by the firms before they choose prices.

To provide a foundation for this reinterpretation, it is important to distinguish between ex ante (total) and ex post (variable) costs where the former includes the costs of installing capacity. Accordingly, one distinguishes between ex ante and ex post reaction curves. When the expected cost of installing new capacity is large relative to the marginal cost of producing or selling output, the firms will end up charging the price that clears the market, given the installed capacities, and will thus behave like a Cournot auctioneer (Beckman 1967; Levitan-Shubik, 1972). To understand this point, define viable capacities as those which would allow the firm to make a long-run profit as a monopolist (hence non-viable capacities will never be selected in equilibrium). When the cost of installing capacity is large, viable capacities necessarily lie under the ex post reaction curves. Suppose for simplicity that consumers are rationed according to the "efficient rationing rule," which says that the highest valuation buyers buy at the lowest price. Then, starting from the market clearing price, it does not pay for any firm to raise its price, that is, to move its quantity away from its ex post reaction curve. Given that firms are already selling their entire capacity, they don't want to lower prices either.

Unfortunately, this straightforward interpretation of Cournot behavior is limited to that class of industries where total investment costs far exceed marginal costs. [20] asks whether the reinterpretation can be generalized to arbitrary investment and production costs. Kreps and Scheinkman assume efficient rationing and show that (if the demand function is concave) the outcome of the two-stage investment-in-capacity and then price-fixing game is indeed the Cournot outcome. While this result does not carry over to alternative rationing schemes (Davidson and Deneckere, 1986), it is of interest. First, the efficient rationing assumption may be a good approximation in some industries; this is in particular the case if there exists a perfect (unrationed) secondary market in which output is resold to the highest bidders. Second, what matters for many applications of the Cournot model is that the reaction curves be downward sloping, not that the reduced-form profit functions have the exact Cournot form. The Kreps-Scheinkman result suggests that such an assumption holds in an interesting class of capacity setting games. Finally, this article was instrumental in popularizing the notion that Cournot quantities might stand for capacity, and more generally technological, choices.

The second paper [23] applies game theory to certain patterns of organizational behavior and offers some insights into the theory of the firm. Pioneering work by Williamson (1975, 1985) and Grossman and Hart (1986) has offered clues as to what determines the governance structure of firms. In particular, transaction-specific investment by parties to a relationship must be protected from expropriation by a long term contract, or, when this is not feasible, by an authority relationship giving the party who invests residual rights of control in contingencies that were not foreseen by the contract. Grossman and Hart trace authority to the ownership pattern. Often, however (and this does not contradict the powerful theory of Grossman and Hart), business relationships are run informally on a day-by-day basis without resorting to long-term contracting or vertical integration, and use reputation to economize on the associated transaction costs (MaCaulay, 1963). As Simon (1951) noted, the parties must have a stake in a maintained relationship in order not to engage into opportunistic behavior when contracts are informal. In [23], the parties' behavior is observable, but not verifiable by a court. That is, a business partner who expropriates the other party's investment does not incur any short-run penalty because no contract is enforceable, but triggers future punishment by either the expropriated business partner or by other potential partners who observe this behavior.[14]

The paper also contains a number of insightful queries. In particular, what is an employee's incentive to maintain a firm's reputation given that the employee's tenure in the firm is shorter than the firm's lifetime? One interesting possibility suggested in the paper is that the employee be transformed into a long-term player through ownership of stock options. The player will then internalize part of the loss of reputation of the firm. This possibility is particularly relevant for high ranking officials as the stock value is more informative about their activity than about the division managers, salespersons, and so on.[15] Another possibility is that the the opportunistic employee faces ostracism from other employees who in turn would face ostracism if they were more accommodating (Cremer, 1986). A third possibility is that the employee has a career concern à la Holmstrom (1982) and does not want to develop a reputation for behaving unfairly. Kreps [23] (see also Simon, 1951; Williamson, 1975) also argues that, for reputation phenomena to work well, the meaning of "appropriate or equitable fulfillment" of a contract must be commonly understood by all. That is, the parties must grasp what outcome and behavior would result

[14]The model in [23] is an infinitely repeated game with complete information about preferences. See Hart and Holmstrom (1987) for a gang-of-four reputation model in this spirit, and Fudenberg, Kreps, and Maskin [35] for a "folk theorem" with opponents that change over time.

[15]In a related contribution, Eaton (1986) shows how finitely-lived individuals may establish infinitely-lived corporate entities to act as banking intermediaries. In his model, individual lenders anticipate lending only once and therefore have no incentive to maintain a reputation to punish default by borrowers. However, they invest their savings in bank deposits and the owners of equity in banks have an incentive to enforce repayment to maintain the value of their equity.

from the (missing) optimal complete contract. [23] interprets this common understanding as (part of) corporate culture.

Foundations of Asset Pricing

A substantial fraction of Kreps's early work studies asset pricing. I will first review his work on dynamically complete markets and then tackle his papers about the martingale approach to asset pricing.[16]

As is well known, the Arrow-Debreu model of complete markets presumes that there are as many markets as there are goods, where goods are defined by physical attributes, location, date of availability, state of nature, and so on. It might thus seem that obtaining Pareto efficient outcomes would require a huge number of markets. However, Arrow (1964) showed that one may be able to save somewhat on the number of markets by using the temporal structure of the economy. One method in an economy with multiple physical goods is to restrict intertemporal trade to a numeraire and set up spot markets to allow trade of goods within a period. Another method, and the focus of [14], is to use the intertemporal resolution of uncertainty.

Suppose, for example, that there is a single physical good, that there are three time periods ($t = 0, 1, 2$) and six states of nature ($\omega_1, \ldots, \omega_6$). At date one, let us say that traders are able to narrow the field of which state of nature will occur at date 2; they know that it will either be in group/event $E_1 = (\omega_1, \omega_2, \omega_3)$ or in group/event $E_2 = (\omega_4, \omega_5, \omega_6)$.

In an Arrow-Debreu economy, the traders can buy or sell at date 0 securities that deliver at date 2 one unit of the good in the state of nature actually observed and 0 in the other states. It might therefore seem that a complete set of markets for consumption at date 2 would have six elements, corresponding to each of the states of nature. But Arrow pointed out that only five markets are needed to obtain de facto complete markets: let the traders have two markets at date 0 to trade securities contingent on which group is observed, and then three markets at time 1 contingent on the state of nature actually observed in time 2. Arrow's idea that one can obtain essentially complete markets by setting up a set of contingent claims that follow the resolution of uncertainty (if traders have correct expectations about future prices) was extended by Guesnerie and Jaffray (1974) with the help of Radner's (1972) framework to settings with more than two periods.

In practice, agents trade assets like stocks and bonds that have yields in many states of nature, rather than Arrow-Debreu securities. However, this does not prevent markets from being essentially complete as long as there are

[16]See Huang and Litzenberger (1988, chapters 7 and 8) for a very clear exposition of the topics covered in this section.

enough linearly independent assets to recreate fictitious Arrow-Debreu securities. The classic example is the case of a binomial resolution of uncertainty in an economy with a stock and a bond. The return on the stock depends on the state of nature while the return on the bond isn't contingent on which state occurs. As an example, suppose that there are four dates and two assets, a stock and a bond, traded at each date. Suppose further that there are eight states of nature and that information accrues in a binary fashion: at date 0, all eight are possible; at date 1, the field is narrowed to four; at date 2, to two; at date 3, to one. The two securities traded three times (which means six markets) allow agents to choose any date-3 state-contingent consumption by varying the proportion of stocks and bonds in their portfolio (as long as the stock and the bond are independent; that is, span today the space of trades contingent on tomorrow's two subevents, which will generally be the case). Thus a stock and a bond suffice to yield essentially complete markets; as a consequence, other claims (such as options written on the stock) are redundant, and therefore priced by arbitrage independently of the investors' preferences.

Kreps [14] uses an elegant and general formalization of Arrow's framework. Suppose that there are T periods at which agents consume and n long-lived assets, where an asset is defined by a claim for the good in each state of nature at the final date T. In each time period t, traders know that the state of nature at that time belongs to some group of state E_t, taken from the set of all possible states of nature possible at the ending date T. Under what conditions can one recreate a complete set of Arrow-Debreu securities?

As it turns out, there should be at least as many independent assets in any time period t as there are groups or events E_{t+1} in the next time period conditional on the group E_t in this time period. [14] goes on to show that if there are at least as many assets as there are possible subevents from one period to the next (that is, as there are E_{t+1} consistent with E_t), then for "almost any" choice of assets (keeping their number and the other features of the economy constant), markets are essentially complete. The only exception would be if some of the assets are not independent of each other, so combining them does not allow full spanning. This result is the logical conclusion of one of the lines of research initiated by Arrow (1964).

Finally, [14] analyzes continuous-time trading of a stock and a bond when uncertainty resolves as a Brownian motion. Black and Scholes (1973) showed that options on the stock are priced by arbitrage; that is, their price depends on the preferences of consumers only through the prices of the stock and the bond.[17] But Black and Scholes also make the extreme assumption that trading

[17] Preliminary intuition about why these two assets suffice to obtain complete markets in spite of the continuum of states of nature is as follows. A one-dimensional Brownian motion is the limit of the binomial resolution of uncertainty paradigm mentioned above, where the convergence is both in distribution (a weak criterion) and in conditional distribution (so that the information structure at each period is approximately the same). And we noted that two independent assets suffice to yield complete markets in the binomial case. This is only an intuition. Furthermore, things are more complex in the case of a multi-dimensional Brownian motion.

happens in continuous time, and one would like to know whether approximately efficient allocations can be obtained with a large but finite number of trading dates for the two assets. The analysis builds on Harrison and Kreps [12] to prove approximate efficiency under some conditions.

Harrison and Kreps [12] contains a fundamental investigation of viable price systems for securities; that is, prices that do not give rise to arbitrage or free lunches. It unifies two different literatures in intertemporal pricing: by extending Ross (1976) to the multiperiod context[18] and by encompassing Black-Scholes (1973) within the general pattern of Ross (1976).

In [12], the set of existing securities may or may not yield essentially complete markets. Harrison and Kreps show that a) given some consumption patterns for the traders and price processes for the assets, there exists a fictitious probability measure on states of nature such that each asset's price (normalized by the price of a bond) follows a "martingale" and b) markets are essentially complete if and only if there exists a unique such fictitious measure.

Result a means that replacing the real probability measure by a fictitious one, the price of an asset is given by the familiar formula of a world with risk neutral investors: today's price is equal to the expectation of tomorrow's price plus dividend. The intuition for this is simple. With risk averse traders, the price of an asset today is equal to the expectation of its price plus dividend tomorrow times the ratio of marginal utilities of consumption tomorrow and today, which expresses the fact that a higher return tomorrow is more valuable when consumption is low then. In equilibrium, this equality must hold for all traders. To obtain a martingale formula, it suffices to define the fictitious probability of a state of nature as its real probability times the ratio of marginal utilities of consumption.[19] As the learned reader will notice, this construction amounts to creating Arrow-Debreu securities and using their values to price assets in terms of their returns in each state of nature. To compute the fictitious probabilities in practice, one writes the martingale equations and inverts them to obtain the probabilities. This may yield a single solution as in the case of essentially complete markets (result b) or multiple ones (indeed a continuum of them) if markets are incomplete.

These results show that one can identify complete markets by demonstrating existence of a unique martingale measure. In addition, constructing the martingale measure considerably simplifies the checking for the absence of arbitrage possibilities, as one does not need to consider all trading possibilities; indeed Wall Street has made substantial use of these fictitious probability measures. The construction of the fictitious measure also simplifies the pricing of options when markets are complete. Finally, the use of fictitious measures

[18]For more on the equivalence between the absence of arbitrage and the existence of linear pricing rule, see Hart (1974, p. 96), Cox and Ross (1976), and Ross (1978).

[19]Two details. First, the ratio of marginal utilities is taken between the initial and final dates. Second, one divides the fictitious probabilities by the initial price of the bond to get a martingale for asset prices normalized by the price of the bond.

simplifies the analysis of a trader's optimal consumption-portfolio policies (Cox and Huang, 1986).

I should also briefly mention Kreps's less foundational work in asset pricing [1, 6, 24]. For example, [1] is a series of insightful observations about the concept of rational expectations equilibrium.[20] It supplies a simple and illuminating example of a future market in which an equilibrium does not exist (see also Green, 1977). The issue is that when the state contingent prices change slightly, the information extracted by the traders from the equilibrium price may change discontinuously. If the price of an asset is the same in two different states, then traders have no way of telling the two states apart. But if the price depends on the state, even by a very small amount, information is available to traders. (As Kreps notes, this is reminiscent of Hart's (1975) example of nonexistence of a competitive equilibrium in an economy with incomplete markets where a small change in asset prices may change the dimension of the consumers' feasible consumption set.) The paper also generalizes a result Kreps attributes to Stiglitz (1971). If there is no insurance motive in a futures market and there are no liquidity traders, then no trader has an incentive to trade in equilibrium even if traders have differential information. Milgrom and Stokey (1982) provide a more general no-trade theorem.

Epilogue

I would like to conclude on a personal note by expressing admiration for Kreps's work, an admiration apparently shared by the entire profession. It goes without saying that he has had much influence on the many researchers whose interests are closely related to his. But while my own research has had little overlap with David's, I have very often borrowed insight and inspiration from his work; the same holds for many other economists not directly engaged in his areas of research. His contributions combine technical virtuosity and economic relevance and have lasting power. The John Bates Clark Medal committee can only be congratulated for its superb choice.

■ *The author is grateful to Chi-fu Huang, David Kreps, Carl Shapiro, Joseph Stiglitz, Timothy Taylor and Philippe Weil for helpful comments, and to the National Science Foundation for financial support.*

[20]This paper has been a source of pain to graduate students at Stanford (at least). It began as a term paper prepared when Kreps was a student, for a course taught by Joe Stiglitz. And it is reliably reported that since then, when asked by students what sort of term paper he would like to receive, Stiglitz indicates that this was "not too bad."

References

Arrow, Kenneth, "The Role of Securities in the Optimal Allocation of Risk Bearing," *Review of Economic Studies,* April 1964, *31,* 91–96.

Banks, J., and J. Sobel, "Equilibrium Selection in Signaling Games," *Econometrica,* 1987, *155,* 647–662.

Beckman, M., "Edgeworth-Bertrand Duopoly Revisited." In Henn, R., ed., *Operations Research-Verfahren,* III. Meisenheim: Verlag Anton Hein, 1967.

Bewley, T., "Thoughts on Tests of the Intertemporal Asset Pricing Model," mimeo, Northwestern University, 1982.

Black, F., and M. Scholes, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy,* 1973, *81,* 637–654.

Cho, I. K., and J. Sobel, "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory,* forthcoming.

Constantinides, G., "Habit Formation: A Resolution of the Equity Premium Puzzle," mimeo, Graduate School of Business, University of Chicago, 1988.

Cox, J., and C. Huang, "Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process," *Journal of Economic Theory,* 1986, *49,* 89–139.

Cox, J., and S. Ross, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics,* March 1976, *3,* 145–166.

Cremer, J., "Cooperation in Ongoing Organizations," *Quarterly Journal of Economics,* 1986, *CI,* 33–49.

Davidson, C. and R. Deneckere, "Long-Term Competition in Capacity, Short-Run Competition in Price, and the Cournot Model," *Rand Journal of Economics,* 1986, *17,* 404–415.

Duffie, D., and C. Huang, "Implementing Arrow-Debreu Equilibria by Continuous Trading of Few Long-Lived Securities," *Econometrica,* November 1985, *53,* 1337–1356.

Eaton, Jonathan, "Lending for Costly Enforcement of Repayment and Potential Fraud," *Journal of Banking and Finance,* 1986, *10,* 281–293.

Epstein, L., and S. Zin, "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica,* 1989, *57,* 759–778.

Farmer, R., "RINCE Preferences," *Quarterly Journal of Economics,* 1990, *105,* 43–60.

Fudenberg, D., and D. Levine, "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica,* 1989, *57,* 759–778.

Fudenberg, D., and E. Maskin, "Folk Theorems for Repeated Games with Discounting or with Incomplete Information," *Econometrica,* 1986, *154,* 533–554.

Fudenberg, D., and J. Tirole, "Perfect Bayesian and Sequential Equilibrium," *Journal of Economic Theory,* forthcoming.

Goldman, S., "Flexibility and the Demand for Money," *Journal of Economic Theory,* October 1974, *9,* 202–222.

Green, J., "The Non-Existence of Information Equilibria," *Review of Economic Studies,* October 1977, *44,* 451–463.

Grossman, S., and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review,* June 1980, *70,* 393–408.

Grossman, S., and O. Hart, "The Costs and Benefits of Ownership: A Theory of Lateral and Vertical Integration," *Journal of Political Economy,* 1986, *194,* 691–719.

Guesnerie, R., and J. Y. Jaffray, "Optimality of Equilibrium of Plans, Prices, and Price Expectations." In Dreze, J., ed., *Allocation under Uncertainty.* London: Macmillan, 1974.

Harsanyi, J., "Games with Incomplete Information Played by Bayesian Players," *Management Science,* 1967–68, *114,* 159–182, 320–334, 486–502.

Hart, O., "Essays in the Economics of Uncertainty," Ph.D. dissertation, Princeton University, 1974.

Hart, O., "On the Optimality of Equilibrium when Markets are Incomplete," *Journal of Economic Theory,* December 1975, *11,* 418–443.

Hart, O., and B. Holmstrom, "The Theory of Contracts." In Bewley, T., ed., *Advances in Economic Theory, Fifth World Congress.* Cambridge: Cambridge University Press, 1987.

Henry, C., "Investment Decisions under Uncertainty: The Irreversibility Effect," *American Economic Review,* December 1974, *64,* 1006–1012.

Holmstrom, B., "Managerial Incentive Problems: A Dynamic Perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck.* Helsinki: Swedish School of Economics, 1982.

Huang, C., and R. Litzenberger, *Foundations for Financial Economics.* Amsterdam: North-Holland, 1988.

Kohlberg, E., and J. F. Mertens, "On the Strategic Stability of Equilibria," *Econometrica,* 1986, *54,* 1003–1038.

Levitan, R., and M. Shubik, "Price Duopoly and Capacity Constraints," *International Economic Review,* February 1972, *13,* 111–122.

MaCaulay, S., "Non-Contractual Relations in Business," *American Sociological Review,* 1963, *28,* 55–70.

Milgrom, P., and J. Roberts, "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory,* August 1982a, *27,* 280–312.

Milgrom, P., and J. Roberts, "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis," *Econometrica,* March 1982b, *50,* 443–460.

Milgrom, P., and N. Stokey, "Information, Trade and Common Knowledge," *Journal of Economic Theory,* February 1982, *26,* 17–27.

Nash, J., "The Bargaining Problem," *Econometrica,* April 1950, *118,* 155–162.

Radner, Roy, "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica,* March 1972, *40,* 289–303.

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

Ross, S., "A Simple Approach to the Valuation of Risky Streams," *Journal of Business,* 1978, *51,* 453–475.

Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit," *Zeitschrift für die gesamte Staatswissenschaft,* 1965 *112,* 301–324.

Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory,* 1975, *14,* 25–55.

Simon, H., "A Formal Theory of the Employment Relationship," *Econometrica,* July 1951, *19,* 293–305.

Spence, M., *Market Signaling.* Cambridge: Harvard University Press, 1974.

Stiglitz, Joseph E., "Information and Capital Markets," mimeo, 1971.

Weil, P., "The Equity Premium Puzzle and Risk-Free Rate Puzzle," *Journal of Monetary Economics,* mimeo, Harvard University 1989.

Weil, P., "Non-Expected Utility in Macroeconomics," *Quarterly Journal of Economics,* 1990, *105,* 29–42.

Williamson, Oliver, *Markets and Hierarchies: Analysis and Antitrust Implications.* New York: Free Press, 1975.

Williamson, Oliver, *The Economic Institutions of Capitalism.* New York: Free Press, 1985.

This article has been cited by:

1. Olivier Mesly, Maria Petrescu, Alexandra Mesly. 2022. Terminology Matters: A Review on the Concept of Economic Predation. *Journal of Economic Issues* **56**:4, 959-987. [Crossref]
2. Manuel Foerster, Achim Voss. 2022. Believe me, I am ignorant, but not biased. *European Economic Review* **149**, 104262. [Crossref]
3. Manuel Foerster, Achim Voß. 2018. Believe Me, I Am Dumb, but Not Corrupt. *SSRN Electronic Journal* **128**. . [Crossref]
4. William E. Becker. Why go to College? The value of an Investment in Higher Education 91-120. [Crossref]