Journal of Economic Perspectives—Volume 31, Number 2—Spring 2017—Pages 237–256

# Yuliy Sannikov: Winner of the 2016 Clark Medal

Susan Athey and Andrzej Skrzypacz

Yuliy Sannikov is an extraordinary theorist who has developed methods that offer new insights in analyzing problems that had seemed well-studied and familiar: for example, decisions that might bring about cooperation and/or defection in a repeated-play prisoner's dilemma game, or that affect the balance of incentives and opportunism in a principal–agent relationship. His work has broken new ground in methodology, often through the application of stochastic calculus methods. The stochastic element means that his work naturally captures situations in which there is a random chance that monitoring, communication, or signaling between players is imperfect. Using calculus in the context of continuous-time games allows him to overcome tractability problems that had long hindered research in a number of areas. Previous models often abstracted from crucial economic forces in the name of tractability, but Sannikov's methods allow models to include the most important forces and thus deliver results that are much more relevant and intuitive. Sannikov's remarkable research agenda has substantially altered the toolbox available for studying dynamic games.

His early training focused on mathematics. In high school, Sannikov won three gold medals at the International Mathematical Olympiads on behalf of his native Ukraine. He studied mathematics at Princeton as an undergraduate, where he was influenced by economist Dilip Abreu and mathematician Yakov Sinai. He completed his PhD in economics at the Stanford Graduate School of Business, where he was

■ *Susan Athey is Economics of Technology Professor and Andrzej Skrzypacz is the Theodore J. Kreps Professor of Economics, both at the Stanford Graduate School of Business, Stanford, California.Their email addresses are athey@stanford.edu and skrz@stanford.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.31.2.237 doi=10.1257/jep.31.2.237

**Yuliy Sannikov**

advised by Robert Wilson and Andrzej Skrzypacz. Beyond his immediate advisors, a number of scholars inspired and contributed to his thinking, including Michael Harrison, Peter DeMarzo, Thomas Sargent, Darrell Duffie, and Paul Milgrom.

The excellence of Sanikov's work is widely recognized. He received the John Bates Clark medal in 2016, which is awarded annually by the American Economic Association "to that American economist under the age of forty who is judged to have made the most significant contribution to economic thought and knowledge." In 2015, the American Finance Association awarded the Fischer Black Prize to Sannikov, honoring an individual researcher under age 40 "for a body of work that best exemplifies the Fischer Black hallmark of developing original research that is relevant to finance practice."

Sannikov's research introduces not only technical advances, but also qualitatively new ideas, which is perhaps the greatest accolade one can bestow on any researcher. A hallmark of his papers is that they take the existing literature to a new level, opening up new lines of inquiry, and allowing qualitatively different types of insights to be derived. This essay offers an overview of Sannikov's research in several areas. We begin with his work using continuous-time approaches, rather than the better-known discrete-time approaches, in the analysis of dynamic games and dynamic contracting. We also sketch some of his more recent work that tackled more complex models in the design of securities, market microstructure, and the role of financial crises in macroeconomics, where the greater tractability available through Sannikov's approaches has an enormous impact in rigorous theoretical

*Table 1*
**Selected Research Papers by Yuliy Sannikov**

| | |
|---|---|
| 1. | "Games with Imperfectly Observable Actions in Continuous Time." 2007. *Econometrica* 75(5): 1285–1329. |
| 2. | "A Continuous-Time Version of the Principal–Agent Problem," 2008. *Review of Economic Studies* 75(3): 957–84. |
| 3. | "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," (with Peter M. DeMarzo). 2006. *Journal of Finance* 61(6): 2681–2724. |
| 4. | "Learning, Termination and Payout Policy in Dynamic Incentive Contracts," (with Peter M. DeMarzo). Forthcoming. *Review of Economic Studies.* |
| 5. | "Moral Hazard and Long-Run Incentives." 2014. Working paper 3430, Stanford Graduate School of Business, https://www.gsb.stanford.edu/faculty-research/working-papers/moral-hazard-long-run-incentives. |
| 6. | "Reputation in Continuous-Time Games," (with Eduardo Faingold). 2011. *Econometrica* 79(3): 773–876. |
| 7. | "Impossibility of Collusion under Imperfect Monitoring with Flexible Production," (with Andrzej Skrzypacz). 2007. *American Economic Review* 97(5): 1794–1823. |
| 8. | "The Role of Information in Games with Frequent Actions," (with Andrzej Skrzypacz). 2010. *Econometrica* 78(3): 847–82. |
| 9. | "A Macroeconomic Model with a Financial Sector," (with Markus K. Brunnermeier). 2014. *American Economic Review* 104(2): 379–421. |
| 10. | "The I Theory of Money," (with Markus Brunnermeier). 2014. Working paper 3431, Stanford Graduate School of Business. https://www.gsb.stanford.edu/faculty-research/working-papers/i-theory-money. |

analysis of some difficult and long-standing problems. Table 1 lists a selection of Sannikov's research papers. In this essay, we will refer to Sannikov's papers by their number in the table, while referring to other papers using the familiar author-date method.

## Overview of Sannikov's Introduction of Continuous-Time Methodology into Dynamic Games

Game theory has traditionally been done (mostly) with discrete-time models, rather than models in which agents make choices continuously. Discrete time makes it more straightforward to define the set of admissible strategies so they map unambiguously to payoffs, and to apply definitions of perfect equilibria. For example, suppose Alice and Bob play a repeated version of the well-known Battle of the Sexes game. In this game, two players are trying to meet at either the opera or the hockey game. Although Alice prefers the opera and Bob prefers the hockey game, they both prefer to be at the same event rather than ending up at different places. Suppose they follow simple strategies of rotating where they meet: going to the opera in even periods and hockey games in odd periods. This is a well-defined pair of strategies in discrete time with a clear outcome. But in continuous time, there are no longer odd and even periods. Moreover, it is not immediately clear how to define this kind of game in continuous time to allow for infinitely frequent changes of actions as a function of time and (even more problematic than in this example), as a function of opponent's past actions.

Or in a different game (one that is known as the Rubinstein bargaining game), suppose that Chris and Dale bargain to split a dollar. Chris makes all the offers and Dale can accept or reject offers. Rejections extend the game. Players prefer early agreements because delay implies some costs (that is, players discount the future relative to the present). In discrete time, if Dale expects a positive amount in the best equilibrium for him, Chris would offer slightly less in the first period and because of discounting, Dale would accept that first offer, making Dale's expectation irrational. Therefore, the unique subgame perfect Nash equilibrium is that Chris offers to take the whole dollar and Dale accepts (Dale is indifferent to accept that offer since he expects never to be offered any amount). But in continuous time, it appears that any split of the dollar is an equilibrium: if Dale expects Chris to offer 80 percent after every history of the game, it is sequentially optimal for Dale to reject any lesser offer—because in continuous time the future arrives without any costly delay.

Of course, there have been some previous applications of continuous-time methods in strategic environments. For example, the "war of attrition" is a game in which players would benefit from outwaiting an opponent, but in which waiting is also costly, and so players need to select a time to stop waiting. Another exception involves differential games, in which the game itself evolves over time. A classic example is the "homicidal chauffeur" problem, in which the driver of a car tries to chase down a jogger. The car is faster, but the jogger is more maneuverable, which sets the conditions for the dynamics of the game to evolve. In the area of contracting, the seminal Holmstrom and Milgrom (1987) paper on "aggregation and linearity" used a continuous-time model to provide a rationale for simple linear contracts in an environment where agents can choose their effort as a function of time, and because of the continuous-time assumption, are able to observe the results of their efforts immediately, not after a delay.

With a few such exceptions duly noted, it's fair to say that before Sannikov's work, the use of continuous-time models in game theory applications was highly limited.[1] A widely recognized problem with this approach was that in repeated interactions where actions are perfectly observed, writing the model in continuous (or almost continuous) time typically creates a situation in which agents who deviate from an equilibrium path in a dynamic game can be immediately punished, which

[1] Outside of the game theory context, continuous-time models had of course been used extensively in several areas of economics, for example in the areas of asset pricing and real option theory. Most of these applications have been in the area of competitive markets or single-person decision problems, rather than strategic interactions.

eliminates any possible benefits of such deviation, and drastically reduces usefulness of the continuous-time formulation. Other related problems with continuous-time modeling are described in Simon and Stinchcombe (1989). Many researchers have proposed fixes to these problems, but it is fair to say that—until the arrival of Sannikov's work—none of these approaches made a broad impact on work in applied theory.

The key early insight in Sannikov's work was that if we introduce imperfect monitoring, a feature of many real-world environments, it becomes possible to define strategies (or contracts) as a function of the imperfect signals that agents observe. If an agent deviates, then in a situation of imperfect monitoring, it takes time for other agents to become confident of the deviation. With this assumption, opponents no longer can react instantly to a deviation by an agent. By making the models more realistic for many applications, Sannikov managed to also achieve new tractability, which then allowed him and other researchers to provide new robust economic intuition.

Sannikov's [2] first paper on strategic interactions in continuous time started when he was a second-year PhD student. He had been studying dynamic contracting models in Thomas Sargent's class, and was presenting a new analysis of the problem studied by Phelan and Townsend (1991), which was an analysis of optimal dynamic contracting with moral hazard and risk-averse agents. The weakness of that paper was that while the result was very general, analytical characterizations were not available; instead, the optimal contract needed to be computed, and it was hard to deliver economic intuition for some features of the computed solution. More specifically, the setting did not rely on the simplifications of the Holmstrom and Milgrom (1987) approach mentioned earlier, and thus the optimal contracts were much more complicated than Holmstrom and Milgrom's linear contracts. Sannikov's insight was that if we take the continuous-time limit of the optimality conditions in Phelan and Townsend (1991), and assume that the noise in output can be described as Brownian motion, the problem becomes drastically simpler.

The key element of the intuition is that in static problems where agents need to be motivated to take an action through incentive contracts, the optimal contract calls for large rewards or punishments when unusual outputs are realized. Indeed, these models predict that contracts should be highly nonlinear, a prediction that is not borne out in many real-world settings. The optimality of large rewards and punishments also arises in a discrete-time formulation of a dynamic game: signals that are very informative about an agent's action call for large changes in continuation payoffs ("continuation payoffs" are the expected present value of equilibrium payoffs in the continuation of the game). As a result, the trade-off between incentives and risk becomes convex. The optimal resolution of this trade-off at a given point of time for a given continuation payoff depends on the principal's entire value function over all continuation payoffs.

In contrast, if information about an agent's performance arrives continuously, as in Holmstrom and Milgrom (1987) and [2], then the way in which the optimal contract changes with the realization of signals today leads to only small changes in payoffs in the future (otherwise, the risk of sharp decreases in future payoffs would expose the agent to a level of risk that outweighs the benefits in terms of incentives). As a result, the trade-off between incentives and risk is local: it depends only on how the value function changes locally, and the optimal contract/set of achievable payoffs can be computed much more easily, as a solution to a differential equation. Unlike in discrete time, where a solution requires knowing what is sequentially credible for large deviations, we only need to know what is feasible locally. Sannikov demonstrated in his work how this insight can be used to provide an alternative analysis of strategic dynamic interactions, one that leads to a more unified understanding of optimal contracts. The focus on local interactions, in turn, leads to a better understanding of long-run properties of the optimal contract through stochastic differential equations (that is, a differential equation that involves both a deterministic and a random element), which are a staple of his analysis.

After demonstrating the value of using continuous-time methods to analyze limits of previously more complicated problems, Sannikov developed formalism to state a general class of problems directly in continuous time, so that a wide range of applications can directly rely on his general framework.

## Foundations of Repeated Games and Economics of Relationships

Repeated games are important in the social sciences because they are the basic model we use to discuss a variety of repeated strategic interactions in which agents face tension between opportunistic and cooperative behavior. Applications range from social dilemmas to team production to tacit collusion. In a series of papers, Sannikov has advanced our understanding of equilibria in repeated games and through that work provided new insights on the economics of ongoing relationships.

Sannikov's dissertation [1] was a technical breakthrough, applying the mathematical tools of stochastic calculus to analyze repeated games. At a time when the vast majority of the literature used discrete-time models, the paper showed how continuous-time tools allowed otherwise intractable games to be analyzed elegantly and neatly. The paper characterizes the set of equilibrium values obtainable in perfect public equilibria of a repeated game with imperfect public monitoring. The key simplifying assumption is that agents choose actions in continuous time and observe public signals about opponent play via a continuous information process. The monitoring technology is modeled as a pair (one for each player) of Brownian motions with a drift controlled by the instantaneous actions of agents.

In this setup, Sannikov showed, it is possible to characterize the boundary of the set of equilibrium payoffs as the solution to an ordinary differential equation. His characterization applies for any fixed discount factor, although the situation in which agents are somewhat impatient is more challenging and more interesting than the analysis of what happens in the limit as players become more patient. In the limit, as players in these games get very patient, tradeoffs between efficiency and incentives often disappear, and all feasible payoffs can be obtained (that is, a "folk theorem" holds). In contrast, with less-patient players, real tradeoffs arise, and a characterization of the set of equilibrium payoffs highlights those tradeoffs. This work was immediately recognized as being path-breaking, a truly significant advance in a literature that had for years been characterized by incremental improvements.[2] In the class of games that Sannikov analyzes, the ordinary differential equation that describes the boundary of the set allows direct interpretation of how dynamic rewards and punishments are used to provide incentives. As Sannikov shows, in his class of games, in the equilibria with payoffs on the boundary of the achievable set, players always remain on that boundary and move continuously in response to the public news.

This technical result has applied consequences. For example, it allows us to understand better the costs to a cartel of not using direct monetary payments between its members and instead using future rewards and punishments—for example, in the form of changes in future market shares. In the optimal collusive equilibrium, if the cartel members adjust their actions frequently and monitoring is via a gradual information process, incentives are provided by small transfers of continuation payoffs in response to the observed realizations of the public signals. The more sensitive are continuation payoffs to these signals, the more high-powered are the incentives faced by firms. The higher the noise in monitoring is, the lower the power of these incentives. When incentives are provided via future reallocation of market shares, typically it reduces somewhat the future total surplus of the cartel. For example, in a repeated prisoners' dilemma, to transfer a future payoff from player 2 to player 1, we need to sometimes let player 1 defect while player 2 is cooperating, which reduces the sum of payoffs over time. This cost of reduced payoffs over time could be avoided if the cartel members engage directly in monetary transfers, for example via product purchases from each other. Harrington (2006) notes that this is commonly done by explicit cartels. Since explicit transfers have their own costs and risks for the cartel (for example, they can increase the risk of detection), for small deviations from allocated market shares, it may be optimal to provide incentives via changes in promised continuation profits (for example, letting firms catch up with their promised sales). This could explain why cartels only meet and settle-up infrequently after unusually asymmetric outcomes.[3]

[2] The previously known method for finding the set of (perfect public) equilibria of such games has been the famous "APS method," named for the Abreu, Pearce, and Stacchetti (1990). This solution method involves iterating on a set of candidate equilibrium payoffs and finding the largest fixed point of the "APS" operator (where this operator takes as input a candidate set of payoffs that can be attained in the future, and returns the set of payoffs that can be obtained if the candidate set is available in future states). The APS method has been widely used in applied work. Despite its success, because the set of equilibrium payoffs is described as a fixed point of an operator, it is often hard to provide analytical characterizations of that set or to say much about the optimal strategies supporting the equilibrium.

[3] Geometrically, this result corresponds to the set of equilibrium payoffs being strictly convex. As Sannikov shows, in the optimal equilibrium there is a one-to-one mapping between the curvature of the set and the cost of using high-power incentives: with curvature, local movements on the boundary create a downward net drift proportional to the sensitivity of payoffs to the observed outcomes. The more curved is the set of continuation payoffs, the higher the incentives to engage in direct transfers. Since curvature is small around the symmetric payoffs, monetary transfers become particularly useful after unbalanced histories that push continuation payoffs away from symmetry.

The property that incremental information affects incentives via transfers of continuation payoffs does not always hold in repeated games in discrete time and/or with other monitoring structures. In some dynamic relationships, it is optimal to react to news by radically changing continuation play. Sannikov's work shows that if actions are frequent and information arrives in small increments, the play in any optimal equilibrium evolves continuously and it is never optimal to use the threat of immediate "price wars" to sustain collusion.

This economic insight helps reconcile theory and empirical evidence. On the one hand, many real-life cartels provide incentives via transfers either through direct transfers or trading of market shares (for example, see Harrington 2006), rather than via the famous price-war mechanism (Green and Porter 1984). On the other hand, when we teach about the economics of relationships by describing cooperative equilibria in repeated prisoners' dilemmas, we introduce the "grim trigger strategy"—if one player defects once, the other player defects forever—as a way of ensuring ongoing cooperation. But most cartels (and other cooperative efforts) clearly do not operate with a grim trigger strategy where one defection ends cooperation for all time. Sannikov's work identifies an important class of situations where players act frequently and information comes continuously for which price wars do not work, because they do not provide incentives without excessively destroying value.

Why might price wars not be useful in providing incentives? In [1], that conclusion emerges magically from continuous-time modeling and the martingale representation theorem, but we may wish to see what economic intuition emerges from a concrete discrete-time setting. This is developed in Sannikov's papers [7] and [8], coauthored with Andrzej Skrzypacz. Those papers study repeated games in discrete time (using standard tools) and then take the limit of those games as agents move closer to continuous time by adjusting their actions more and more frequently. They characterize how continuous-time information can and cannot be used to provide dynamic incentives.[4]

As an example, consider two symmetric firms repeatedly choosing quantities (or team members providing effort) and suppose that market prices depend on the sum of quantities and noise. Thus, prices affect the firms' profits and they also serve as signals about otherwise unobserved quantities produced by other firms. Firms can use prices to test the hypothesis that they are maintaining collusive quantities $(q, q)$ against an alternative that one of the firms deviates and chooses quantity $q' > q$ instead. When noise takes the form of a Brownian motion, the key property of information flow from prices is that the likelihood ratio for any such test is evolving continuously. Thus, in settings of [7] and [8], information flow is proportional to the length of the discrete time period $\Delta$ between points at which players have the opportunity to adjust their actions. Longer periods of observation provide more accurate signals. This is the

[4] The exercise is similar to that of Abreu, Milgrom, and Pearce (1991), but while that paper considered only strongly symmetric equilibria with Poisson information structures, Sannikov's work considers all pure-strategy perfect public equilibria and a mixture of Poisson and Brownian news (in [8]).

kind of monitoring technology we would get if we aggregated information as in the continuous-time formulation in [1] over periods with a finite length.

Suppose firms wanted to hold down quantities produced and collude, but they know that individual firms will be tempted to deviate secretly and to produce higher quantities. If there were no noise in prices, it would be an equilibrium strategy to maintain high prices by the grim trigger strategy: produce at the lower collusive level as long as you see no deviation, but start a price war once you see evidence that any other producer is deviating. But with noise, deviations are not directly observable. Instead, firms have to use prices to conduct statistical inference of when to start a punishment phase: for example, the solution of Green and Porter (1984) is to trigger the price war when prices happen to be unusually low.

When firms aggregate information over a discrete-time period of length $\Delta$ to decide whether to trigger a price war, two quantities are crucial to the effectiveness of collusion. The first is the *likelihood difference* of the test, which at an intuitive level is the amount by which a deviation would increase the probability of the price war. The likelihood difference is key to the strength of incentives. The second is the probability of type-I errors, which in this case is the risk of triggering a price war even when nobody has actually deviated, and the price was low only because of randomness in demand. These errors are the costs of providing incentives via price wars; they quantify the benefits of collusion that are sacrificed per period to provide incentives.

When information flow is continuous, as in [7], any test with a likelihood difference of order $\Delta$ (the length of the period) must have a large probability of type-I error when $\Delta$ is small (about of order $\Delta^{1/2}$). Thus, if collusion brings a gain in payoffs on the order of $\Delta$ per time period (of length $\Delta$) relative to static Nash, the cost of providing incentives is on the order of $\Delta^{1/2}$. Since costs outweigh the benefits when $\Delta$ is small, players are not able to obtain payoffs higher than the static Nash equilibrium. This observation motivated the "Impossibility of Collusion" in the title of [7]. Of course, the takeaway from this paper is not that collusion is impossible, but why many cartels work by transferring payoffs, often at the increased risk of detection, rather than using price wars. The key to determining whether the threat of price wars can be effective in providing incentives depends on the probability of type-I errors compared to the length of the period.

To sum up, there is a tension between the discrete-time theory of collusion, following the work of Green and Porter (1984) and Abreu, Pearce, and Stacchetti (1986), which focuses on the role of price wars in sustaining collusion, and empirical evidence, like that in Harrington (2006), which shows that many cartels prefer to use transfers and/or trade market shares. This tension is reconciled in [7] and [8] by showing that threats of sudden price wars are ineffective when information arrives continuously and players can react to it quickly. This same observation arises in the characterization of equilibrium payoffs of [1] via an ordinary differential equation, which summarizes the cost of incentive provision via transfers of continuation values by the curvature of the equilibrium payoff set. In this characterization, price wars may happen on the equilibrium path, but only as a last resort after the possibilities to transfer payoffs have been exhausted.

These results raise some interesting follow-up questions. First, one may wonder about the Mirrlees critique of principal–agent settings—an observation that when the actions of an agent affect the mean of a normally distributed output variable, it is possible to approach first-best outcomes by imposing extreme punishments on unlikely tail events. This critique has a particularly stark manifestation in the setting of Holmstrom and Milgrom (1987). In the continuous-time model analyzed by that paper, the optimal (second-best) contract is linear—there is a real trade-off between risk and incentives. In contrast, if in the same model, periods instead had length $\Delta$ even for a very small $\Delta$, it would be possible to approach the first-best outcome by highly nonlinear contracts that impose extreme punishments after unlikely tail events. Why does the same observation not apply to repeated games, given that [1] is cast directly in continuous time as in Holmstrom and Milgrom, while [7] and [8] are cast in discrete time, taking the time period $\Delta$ between actions to 0? It turns out that in repeated games, the Mirrlees critique does not apply for several reasons. The transfers are limited by the set of feasible payoffs, so it is not possible to use very large punishments with arbitrarily low type-I errors. More importantly, even if incentives are provided only via transfers of continuation payoffs, as in [1] and [8], it turns out that large transfers will also cause large efficiency losses (because of the curvature of the set of continuation payoffs).

Second, it matters how we take the limit of games towards continuous time—in particular how we model the monitoring technology as $\Delta$ gets smaller. There are many discrete-time processes that converge in some sense to the Brownian motion as $\Delta$ gets small. Yet for monitoring purposes, they have very different properties, and small details can have a big impact on the equilibrium. For an extreme example, we can construct two seemingly similar processes with binary signals in discrete time that converge to the same Brownian motion (with drift that depends on actions). But in the first process, deviations change only the probability distribution over the binary steps, while in the second process they also change the domain of the steps. Games with the first type of process will have very similar properties to the ones described in [1], [7], and [8]. Games with the second process will allow perfect monitoring, where collusion is always easy for small $\Delta$. Fudenberg and Levine (2007, 2009) as well as Sadzik and Stacchetti (2015) provide other examples, in which discrete-time signal processes converge to their continuous-time counterparts, but the informativeness of the signals may or may not converge. Hence, one has to be careful in interpreting games with frequent actions. In [7] and [8], Sannikov creates per-period distributions by integrating continuous-time processes instead of taking limits of discrete-time processes, and this distinction matters for the results.[5]

[5] One can also use this reasoning to compare Poisson and Brownian monitoring technologies. This involves again comparing the *likelihood differences* and *probabilities of type I error* along the lines we described above. Roughly, if deviations of players from the equilibrium path strategies increase the arrival rate of signals (the "bad news" case), *probabilities of type I error* go down to zero at the rate of $\Delta$, so price wars can sometimes help, while if the deviations reduce the arrival rate (the "good news" case), price wars are even less useful than in the Brownian case.

Finally, it is important to ask whether these results have any empirical support. As we mentioned before, there is indirect evidence from real-life cartels that try to collect data on individual behavior of its members and provide incentives via transfers instead of price wars. For example, Cabral (2005) describes early problems of the lysine cartel, which tried to sustain collusion without identifying potential deviators but later introduced an additional system for monitoring market shares. The facts described in these papers are consistent with the main takeaway from this area of Sannikov's research. Those models offer a much better understanding of why in economic relationships that suffer from imperfect monitoring, it is important *not* to use strategies like the grim trigger or price war threats, but to develop more complicated schemes based on trading favors. Several authors have also provided evidence that Sannikov's models may be useful for understanding experimental data (for example, Bigoni, Potters, and Spagnolo 2012), but the literature on games with frequent interactions is still in a relatively early stage. One of the problems is that in discrete-time repeated games with imperfect monitoring, it has been commonly observed that subjects tend to underreact to the introduction of noise, so it is possible that difficulties of experimental subjects in performing statistical inference would swamp the game-theoretic considerations.

## Applications to Dynamic Incentives

Sannikov's models have had very fruitful applications in contract theory. In the application in [2], discussed earlier in this paper, Sannikov analyzes a continuous-time model in which an agent controls the drift of a diffusion process and the paper characterizes the optimal contract in this environment. The modeling is beautiful, and it provides new insights that were not available with existing models. As in traditional models, eliciting higher effort from an agent comes at the cost of exposing the agent to additional risk, so the underlying tradeoff is familiar. However, the dynamics are much richer in this modeling framework. For example, one result is that agents eventually "retire," either because they have had a series of bad luck leading their utility to be so low that it is too expensive to provide additional incentives because the agent can't be hurt any further; or else because they have had a series of good luck leading the utility to be so high that additional rewards are not effective relative to their cost to the principal.

These insights can be taken in a variety of directions. For example, a classic problem in game theory looks at the situation of a "large" (long-lived) individual attempting to establish a reputation with a set of "small" (short-lived) players who react to the large player's behavior. In [6], Sannikov re-examines this problem through the lens of continuous time. The paper is able to characterize equilibria for a range of discount factors. A key qualitative insight from this work is that it is possible to better understand two important forces in models of reputation. The first is the standard "repeated game" force, where the large player is disciplined by the possibility of future punishments. The second is the "reputation" force, where the large player is disciplined by the extent to which today's action affects tomorrow's beliefs. In [6], only the belief force is at play because with Brownian-noise imperfect monitoring, it is not possible to provide incentives via coordinated punishments of the short-lived players for the same reasons as it is hard to sustain collusion with the threat of price wars. Thus, the equilibria will depend only the tradeoff between the reputation of the large player and how it is affected by beliefs. In this setting, the best equilibria turn out to involve mixed strategies that are called "Markov equilibria," in that strategies depend only the current value of the payoff-relevant state variable, rather than the entire history of the game.

Other authors have been building on Sannikov's framework. For example, Bohren (2016) analyzes a market where a firm or worker provides a sequence of customers with a product or service. The worker's effort noisily impacts a payoff-relevant state variable, such as a measure of the worker's rating or a firm's quality. Again, all equilibria are Markov. The model is applied to analyze the optimal design of a platform-based rating application. The paper observes that the modeling "demonstrates the power of the continuous-time setting to deliver sharp insights and a computationally tractable equilibrium characterization in a rich class of dynamic games." It illustrates the power of Sannikov's foundational work.

## Dynamic Contracts, Security Design, and Firm Financing

Continuous-time methods have been central in finance at least since the pioneering work of Black and Scholes (1973) and Merton (1973), but their use has been primarily in the area of asset pricing. Sannikov's methods for analyzing dynamic contracting models extend this technology to corporate finance questions, enabling researchers to analyze the dynamics of optimal executive compensation and security design. As noted at the start, Sannikov received the 2015 Fisher Black Prize in finance in recognition of the importance and usefulness of his contributions.

A pair of papers on finance that Sannikov wrote with Peter DeMarzo [3, 4] exemplify this work. The first paper [3] studies how optimal dynamic contracts can be implemented with a standard capital structure consisting of a credit line, long-term debt, and equity. In the model, a cash-constrained risk-neutral agent undertakes a project that experiences fluctuations in cash flows and thus requires financing by a third party. The agent's effort shifts the mean of cash flows (which can be alternatively interpreted as the agent refraining from diverting cash for private benefit), and the resulting output is observed in real time by the principal/investors. Outside investors are risk-neutral, deep-pocketed, and more patient than the agent. Investors enter into a dynamic contract with the agent in which they commit to payments and a termination rule, both as a function of the project's realized performance. To provide incentives, the optimal contract needs to give the agent some "skin in the game"—which means that total consumption of the agent needs to be sensitive to observed cash flows. This sensitivity is achieved by adjusting the agent's continuation payoff up or down in response to earnings surprises. If performance is good, the agent receives direct payments once the continuation payoff exceeds a threshold level. If performance is bad, the agent is terminated once his continuation payoff falls below a cutoff. Thus, the agent is rewarded for good performance both by accelerating the timing of payouts and reducing the risk of termination. Although termination is inefficient (the project is productive), the possibility of termination is needed to provide incentives because of the agent's limited liability (his continuation payoff cannot be reduced below his outside option). The optimal contract defers payments to the agent to reduce the risk of termination, but must balance this against the agent's relative impatience. The paper first describes the optimal evolution of the agent's continuation payoff and consumption and then describes how the optimal contract can be implemented using standard securities. Despite the complications of dynamic contracting, there exists a relatively simple implementation: the agent gets debt financing to start the project and can draw on a credit line to cover bad outcomes, but only up to a limit—and when the agent hits that limit, the project is terminated. On the other hand, once the agent pays off the credit line, the agent *chooses* to pay dividends to equity holders (which include himself).

While the qualitative implementation of the dynamic incentive contract in terms of a simple capital structure had previously been shown in a discrete-time setting by DeMarzo and Fishman (2007), the continuous-time model with Brownian shocks introduced in [3] provides a major simplification of the problem. It develops a more complete characterization of the optimal contract via an ordinary differential equation, which allows a straightforward calculation of payoffs, security valuations, and comparative statics. For example, the use of credit versus long-term debt varies with features of the environment such as volatility: the less volatile are the project cash flows, the smaller is the credit line given to the agent and hence the smaller room for error the agent has. The paper also shows that the optimal contract may require the firm to hold a compensating cash balance while borrowing (at a higher rate) through the credit line. Thus, the model provides a justification for behavior that might seem irrational absent incentive problems—the cash balance ensures that the firm will have cash flow in future states where investors may not be willing to provide funds.

An additional feature of the implementation is that the contract allows the agent to determine the firm's payout policy and choose which securities to pay off first. The agent chooses to pay off the credit line before paying dividends, but once the credit line is paid off, the agent will pay dividends rather than "hoard cash" (that is, increase the cash balance or pay off long-term debt). The paper shows that volatility affects primarily the mix of securities used by firms—more volatile firms use a larger credit line relative to long-term debt—but has a much smaller impact on the total credit available.

Finally, despite the presence of leverage, the usual conflicts between debt and equity need not arise; that is, neither equity-holders nor the agent have the incentive to increase risk, or to increase dividends to induce default on debt, or to contribute more capital to postpone default. This surprising result arises from the endogenous nature of the optimal contract. For example, because the agent is terminated and his access to funds is cut off once the credit line reaches its limit, the agent has no incentive to take on a high degree of risk. The usual intuition that firms close to bankruptcy have an incentive to take risk that might impose large losses on creditors is overturned because, with continuous Brownian-motion driven cashflows, creditors can stop the agent before the losses become too severe. While this result depends on the continuous nature of the shocks in the model, it reveals that whether asset substitution or excessive risk taking is a first-order problem in firm financing or other dynamic moral hazard problems depends on the nature of the risk: increases in continuous volatility are less important than "tail risk"—the agent taking actions that with small probability can cause dramatic losses. This line of reasoning has been developed for example by DeMarzo, Livdan, and Tchistyi (2014), Biais, Mariotti, Rochet, and Villeneuve (2010), and Varas (2013).

The follow-up study [4] enriches the model of [3] to allow for dynamic learning about the profitability of the project on behalf of both the principal and the agent. Since the agent's effort affects output, which is the source of the principal's learning, the agent has an incentive to manipulate the principal's belief about the project as well as about his effort. In particular, by shirking and reducing output today, the agent obtains an immediate private benefit and lowers the principal's expectations for the output that agent should deliver in the future. This paper develops a rich characterization of how such incentives must be controlled in the optimal contract.

The resulting model produces a very natural "life cycle" of firm dynamics. In its early stages, the firm is financially constrained, with no payouts and the potential for inefficient termination if performance is sufficiently poor. If early-stage termination is avoided, however, the project "matures" and the firm pays dividends (of which the agent receives a share). Notably, dividends are based on expected future earnings and thus are much smoother than the firm's realized earnings. The intuition for this dividend-smoothing is that surprises in current earnings will change beliefs about the firm's future prospects, making it optimal for the firm to absorb the shock to earnings via its cash reserves (for example, a positive surprise raises expectations and makes the project more valuable, so higher reserves are warranted as greater insurance against future uncertainty).

The paper also highlights that features of the optimal contract depend on whether the information the agent can manipulate is project-specific or whether it reflects the agent's general ability (in which case the agent can continue to benefit from this information in his next job, even if he is fired). In the latter case, compensating the agent with equity is sufficient, and once the firm matures, there is no longer any risk of inefficient termination. But in the former case, the paper shows that if the project does poorly in the early stages, even if the firm survives to maturity, the contract is permanently affected—dividend payments are lower, and the agent faces a permanently higher threshold for termination. These long-run distortions are actually optimal because they make it easier to provide incentives (and lower the risk of termination) in the critical early phase of the project.

In [5], Sannikov analyzes optimal contracts in an environment where the agent's effort has a long-run impact on the stochastic process of output rather than just shifting the mean of contemporaneous payoffs. Although this type of setting is obviously much more realistic than typical models where effort only affects current output, dynamic models like this have resisted analysis in the past due to tractability challenges. The paper provides a characterization of the optimal contract, which has some interesting features. First, the agent's exposure to firm risk in the optimal contract is dictated by the degree of control over current outcomes that arises through current as well as past actions. As a result, risk exposure starts small but adjusts towards a target level of risk exposure over time. Second, the contract includes consumption-smoothing features, so that incentive effects of current performance are distributed over time, both on the positive side and on the negative side, to give the most bang-for-the-buck in terms of providing incentives. Third, due to participation constraints (specifically, limited liability), pay-for-performance has bounded sensitivity, and an agent is terminated if performance is too poor.

## Financial Frictions in Macroeconomic Models

Sannikov has been developing a new line of research in macroeconomics with Markus Brunnermeier [9]. This research continues the themes from his previous work, where his modeling approaches are simultaneously more realistic and more tractable than previous models. The inherent complexity of macroeconomic models enables even greater value-added in terms of being able to capture important macroeconomic phenomenon.

After the recent financial crisis, distortions in the financial sector have received renewed attention when modeling fluctuations in aggregate economic activity. The idea that financial distortions play a central role in at least some crisis periods, and that they play a role in the propagation of other types of macroeconomic disturbances as well, is not new. For example, in work that started in the 1980s, Bernanke and Gertler introduced a "financial accelerator" mechanism into dynamic stochastic general-equilibrium models of aggregate fluctuations (for a survey, see Bernanke, Gertler, and Gilchrist 1999). In order for models to be tractable, the models focused on the analysis of fluctuations of the key variable (the net worth of leveraged investors who undertake risky investments) around a constant long-run, steady-state value of aggregate output, and exogenous shocks had to be small enough not to move too far away from this steady state. This in turn allows the dynamics to be represented as approximately linear, which facilitates characterization of equilibrium investor behavior.

Although the tractability implied by linearized representations of equilibrium dynamics have allowed financial accelerator mechanisms to be studied in a variety of contexts, the simplifications also rule out some very important forces and associated phenomena. In particular, the models do not allow the possibility that severe financial crises can occur except as highly unlikely occurrences where exogenous shocks are very large. In addition, the assumptions of the models imply that shocks have the same effect no matter what the underlying state of the economy when the shock occurs, so that the models do not admit the concept of a fragile state of the financial sector of the economy. Thus, it is difficult for these models to capture some of the key stylized facts of the financial crisis.

In [9], Brunnermeier and Sannikov analyze a simple dynamic stochastic general equilibrium model with a financial accelerator mechanism that is in many ways fairly standard. The paper departs from the literature by using continuous-time methods that allow the equilibrium to be characterized by solving an ordinary differential equation without requiring any linearization of the model's dynamics. This allows the dynamics to be characterized under assumptions that do not imply the existence of a "long-run steady state" near which the dynamics are locally mean-reverting. This approach allows the authors to answer an important criticism of the earlier generation of dynamic stochastic general equilibrium models with financial accelerator mechanisms, namely that the models were unable as a quantitative matter to generate large enough aggregate fluctuations in response to aggregate shocks of a realistic size.

Like the previous literature, the key state variable in the model is the net worth of leveraged investors as a share of total wealth. The continuous-time model makes it possible to solve for the amount of time this state variable spends far away from its long-run average value. In turn, this enables analysis of how model parameters affect the frequency and duration of spells where the economy is (endogenously) in a state such that financial distortions are much larger than usual. Financial innovations for risk sharing among individuals, such as derivative contracts that improve risk sharing, or securitization, can create conditions that make financial crises more likely.

The nonlinear solution for the model dynamics also makes it possible to analyze how both average behavior and the way that behavior should respond to further small shocks change depending on the current value of the net-worth state variable. This makes it possible to identify different "phases" of a "financial cycle," and to consider how desirable policy might differ depending on the phase. As a consequence of this aspect of their analysis, the model differentiates between two sources of risk in the economy. In the first, risk is high for exogenous reasons (future disturbances have high variance), while in the second, the risk is endogenous. An example of the second type of risky state is one that emerged because leveraged investors' net worth has declined in response to past shocks. Interestingly, an increase in exogenous risk can lead to a decrease in equilibrium leverage, decreasing the second type of risk. The authors refer to this as the "volatility paradox," and argue that the period of superficially low macroeconomic risk that extended from 1985–2005 helped create conditions that lead to an increase in investor leverage, creating (endogenously) the risk of financial crisis.

In a paper [10], which is not yet published, the authors use their framework to study the effects of monetary policy. In order to do so, the model is augmented with intermediaries that take deposits and make loans. The intermediaries have three functions: 1) monitoring projects; 2) creating a diversified portfolio of projects; and 3) investing in long-term assets and issuing short-term liabilities, thus transforming the maturity of financial instruments. The model establishes that the supply of credit and liquid assets will vary with the underlying state of the economy (similar to the previous model, the state is characterized by the net worth of intermediaries). When borrowers experience shocks and default on loans, banks both reduce lending and supply less "inside money," creating further defaults, so that shocks are amplified by the response of the intermediaries. In this context, monetary policy can affect the probability of financial crises through its effect on leveraged intermediaries. This contrasts with the more traditional mechanism where changes in real interest rates affect household savings and firm investments, and also from existing studies that emphasize the role of bank reserve requirements. The authors discuss the impact of different policy responses in states of the economy characterized by risks due to excess leverage of intermediaries. The authors argue for a policy of "stealth recapitalization" through interest rate cuts, as this type of policy helps strong institutions more than weak ones, because stronger institutions attained their position of strength by hedging their risk. The strong institution carries out the hedge by buying long-term bonds whose value increases when interest rates fall. In contrast, bank bailouts undermine the incentives of banks to manage risk.

While it is early to judge the eventual impact of the arguments in this paper, the approach does bring issues that have been highlighted by the recent financial crisis and the policy response to it to the forefront. The paper further illustrates the fruitfulness of the authors' modeling approach for addressing subtle issues of considerable importance for the applied literature.

## Some Concluding Thoughts

A hallmark of Sannikov's research is that his innovative approaches allow both more realistic assumptions and more realistic conclusions. In many (though certainly not all) settings, imperfect observability of actions or other key economic variables is an empirical reality. It turns out that in such settings, the technical drawbacks of continuous-time models largely disappear. Further, the improved tractability of the setup enables richer, more realistic conclusions as well, and conclusions that do not rely on limiting analysis, like looking only at limits as players become perfectly patient.

One potential criticism of Sannikov's approach is the claim that in reality, strategic interaction never takes place in truly continuous time. The point is fair enough, as long as one also recognizes that discrete-time models of strategic interaction are also a considerable simplification of reality. A benefit of Sannikov's approach is that once analysis is completed, it is possible to find discrete-time analogs of the model and results, and then one can judge if the model is missing something of first-order importance and which insights are likely to be robust. The ability to toggle back and forth between discrete-time and continuous-time models, and to consider the difference between them, is a considerable advance.

Another possible drawback of Sannikov's general approach is that some functional form assumptions are typically required: for example, a common assumption is that the random variation in key economic quantities follows a Brownian motion. However, we believe that in this class of problems, these functional form assumptions are a small price to pay. Economic modeling always involves some simplifications of reality, and the assumptions we make need to be judged on their relative realism and on the power of insights they bring us. Many dynamic agency models become so complicated that they fail to deliver clear intuitions. Sannikov's work has demonstrated that while at first it may seem that continuous-time methods are more complicated than discrete-time methods, after some initial investment they often deliver huge improvements in tractability. His work and the literatures that it is inspiring are allowing researchers to develop new and clarifying intuitions for a variety of problems in which we want to capture agency issues in dynamic settings.

■ *We are grateful to Michael Woodford for his insights surrounding Sannikov's work on financial frictions and to Peter DeMarzo for reviewing a draft of this paper.*

## References

**Abreu, Dilip, Paul Milgrom, and David Pearce.** 1991. "Information and Timing in Repeated Partnerships." *Econometrica: Journal of the Econometric Society* 59(6): 1713–33.

**Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1986. "Optimal Cartel Equilibria with Imperfect Monitoring." *Journal of Economic Theory* 39(1): 251–69.

**Abreu, Dilip, David Pearce, and Ennio Stacchetti.** 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica: Journal of the Econometric Society* 58(5): 1041–63.

**Bernanke, Ben S., Mark Gertler, and Simon Gilchrist.** 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." Chap. 21 in *Handbook of Macroeconomics,* vol. 1, edited by John B. Taylor and Michael Woodford, 1341–93. Elsevier.

**Biais, Bruno, Thomas Mariotti, Jean-Charles Rochet, and Stéphane Villeneuve.** 2010. "Large Risks, Limited Liability, and Dynamic Moral Hazard." *Econometrica* 78(1): 73–118.

**Bigoni, Maria, Johannes (Jan) J. M. Potters, and Giancarlo Spagnolo.** 2012. "Flexibility and Collusion with Imperfect Monitoring." CEPR Discussion Paper DP8877. Available at SSRN: http://ssrn.com/abstract=2034095.

**Black, Fischer, and Myron Scholes.** 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81(3): 637–54.

**Bohren, Aislinn.** 2016. "Using Persistence to Generate Incentives in a Dynamic Moral Hazard Problem." https://ssrn.com/abstract=2889017.

**Cabral, Luís M. B.** 2005. "Collusion Theory: Where to Go Next?" *Journal of Industry, Competition and Trade* 5(3–4): 99–206.

**DeMarzo, Peter M., and Michael J. Fishman.** 2007. "Optimal Long-Term Financial Contracting." *Review of Financial Studies* 20(6): 2079–2128.

**DeMarzo, Peter M., Dmitry Livdan, and Alexei Tchistyi.** 2014. "Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives." Working Paper 3149, Graduate School of Business, Stanford University.

**Fudenberg, Drew, and David K. Levine**. 2007. "Continuous Time Limits of Repeated Games with Imperfect Public Monitoring." *Review of Economic Dynamics* 10(2): 173–92.

**Fudenberg, Drew, and David K. Levine.** 2009. "Repeated Games with Frequent Signals." *Quarterly Journal of Economics* 124(1): 233–65.

**Green, Edward J., and Porter, Robert H.** 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica: Journal of the Econometric Society* 52(1): 87–100.

**Harrington, Joseph E., Jr.** 2006. "How Do Cartels Operate?" In *Foundations and Trends in Microeconomics* 2(1): 1–105.

**Holmstrom, Bengt, and Paul Milgrom.** 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica: Journal of the Econometric Society* 55(2): 303–28.

**Merton, Robert C.** 1973. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science* 4(1): 141–83.

**Phelan, Christopher, and Robert M. Townsend.** 1991. "Computing Multi-Period, Information-Constrained Optima." *Review of Economic Studies* 58(5): 853–81.

**Sadzik, Tomasz, and Ennio Stacchetti.** 2015. "Agency Models with Frequent Actions." *Econometrica* 83(1): 193–237.

**Simon, Leo K., and Maxwell B. Stinchcombe.** 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica* 57(5): 1171–1214.

**Varas, Felipe.** 2013. "Contracting Timely Delivery with Hard to Verify Quality." August 26. Available at SSRN: http://ssrn.com/abstract=2317978.