In this issue

('Foreword' by Eric Maskin)

Game theory

I am delighted to have the opportunity of writing a few words about the eight papers in this special section on game theory. The papers made good reading. Not only do they provide some of the latest results in their respective fields, they also give the reader a good idea of what the next steps forward ought to be.

'Existence of stationary equilibrium for mixtures of discounted stochastic games' by Krishnamurthy *et al.* (**page 1003**) is the latest advance in specifying circumstances in which a stationary equilibrium – an equilibrium in which strategies depend only on the current state – exists in a discounted stochastic game. This research line has a long history, going back to Shapley's fundamental paper of 1953 and includes important previous contributions by Parthasarathy (one of the co-authors of the abovementioned paper). Research has proceeded by successively generalizing the underlying assumptions (zero-sum games were generalized to non-zero games, and finite state-spaces were generalized first to countable and then to uncountable state spaces). From very recent work by J. Levy – who provides counterexamples – we know that stationary equilibria do not always exist in non-zero-sum discounted games with uncountable state spaces. Thus the research agenda is to map out the boundary between existence and nonexistence. The current paper takes an important step in this direction, by establishing existence for mixtures of several classes of games in which existence had previously been shown.

'On Nash-equilibria of approximation-stable games' by Awasthi *et al.* (**page 1014**) examines the problem of computing Nash equilibria. In general, this problem is hard: the time required to find a Nash equilibrium in typical algorithms increases exponentially with the size of the players' strategy sets. The paper shows, however, that for a game in which all 'approximate' Nash equilibria are near a true equilibrium, computation is dramatically simplified: In this case there are algorithms for which running time increases only polynomially. This strikes me as a significant step in the quest to find algorithms that work efficiently in more general classes of games.

Mechanism design theory is the part of game theory that examines how and when we can reach a given economic or social goal by designing a game whose equilibria implement that goal. Of course, which equilibrium concept is adopted will affect which goals can be attained. One particularly attractive equilibrium concept is *strategy-proofness*, wherein each player's equilibrium strategy is optimal for him regardless of the strategy choices by others. Unfortunately, there are many mechanism-design settings for which strategy-proofness is too demanding; there are *no* strategy-proof mechanisms implementing the designer's goals. 'Approximate strategy proofness' by Parkes (**page 1021**) provides a useful review of recent attempts to relax strategy-proofness without giving up its essence altogether. Specifically, the paper explores the possibility that a player: (1) minimizes regret, (2) adopts price-taking behaviour, (3) cannot choose strategies that are too computationally complex, or (4) operates in an environment in which there are many other players, as ways of getting at the simplicity of strategy-proofness while still allowing goals to be implemented.

How to compute equilibria in a multi-market competitive model is a major theoretical and practical problem. Algorithms that find equilbria in polynomial time in Irving Fisher's competitive model have been known for about a decade, but it would be even better to find a way to formulate the Fisher model as a linear program. The paper 'A simplex-like algorithm for linear Fisher markets' by Adsul *et al.* (**page 1033**) takes a significant step in that direction by showing that there exists an algorithm similar to the simplex method for linear programming that can be used to solve a linear version of the Fisher model.

A good many results in the mechanism design literature, which I have already mentioned, are perfectly general – be they positive or negative findings. However, quite a few major propositions depend on agents' types being *one-dimensional*. For example, the revenue-maximizing auctions literature, initiated by R. Myerson, J. Riley and W. Samuelson, assumes that each buyer can be described by a single real number corresponding to his valuation for the good being sold. The paper 'Multidimensional mechanism design: key results and research issues' by Mishra (**page 1043**) is a helpful survey of the (rather limited) literature in which agents' types may be *multidimensional*. It also points out some of the open research questions in the area.

The Gale–Shapley matching algorithm has become the cornerstone of a large theoretical and applied literature that considers everything from matching students with schools to matching kidney donors with kidney recipients. In 'Stable matchings and linear programming', Vohra (**page 1051**) shows that the matching problem can be formulated as a linear program, and that the existence of a stable matching can be proved via a suitable dual algorithm for this program. The paper thereby draws a nice connection between the programming and matching literature.

One important application of mechanism design has been to develop practical methods by which a monopoly seller can sell goods to a set of buyers. However, as McAfee and Vassilvitskii (**page 1056**) point out in 'An overview of practical exchange design', there has been relatively little written about the practicalities of setting up an *exchange* in which there are *multiple* sellers as well as buyers. Their paper sets out some of the basic principles and objectives that should help guide future work in this area.

Finally, Truthful multi-armed bandit mechanisms for multi-slot sponsored search auctions by Das Sharma *et al.* (**page 1064**) examines the problem of designing a mechanism for determining which ads are displayed to a search-engine user and how much the advertisers pay for these ads, on the basis of (1) the keywords the user selects and (2) which ads the user clicks on. The goal is to maximise the advertisers' social welfare, and the paper shows how this can be accomplished through a multiround auction in which advertisers are induced to bid their true willingness to pay, and the early rounds reveal the click probabilities.

I hope the readers of these papers enjoy them as much as I have.