INTRODUCTION

- In this section of the course we are going to study the functioning of competitive markets with many buyers and sellers.
- The principal trading mechanism used in these markets will be the **double auction**:
  - Each buyer may enter two types of orders:
    1. **Market order**: buy a specified amount of the good at the best available offer price. In a simple case of one-unit orders this translates into buying the unit with the lowest offered price. Market orders are *filled instantly* (conditional on sufficient supply).
    2. **Limit order**: by a specified amount of the good at the price that is limited by a specified maximum limit. If there are units that are offered at a lower price, various allocation rules can be used, but usually the difference is split evenly between the buyer and the seller. Limit orders are *filled at the earliest opportunity*, and depending on availability, may only be filled partially. Therefore some limit orders may allow an *all-or-nothing option*. Furthermore, limit orders may be placed for a certain time period, after which they are automatically cancelled, or they may be placed as *good until filled or cancelled*.
  - Each seller may enter two types of orders as well:
    1. **Market order**: sell a specified amount of the good at the best available demand price. In a simple case of one-unit orders this translates into selling the unit to the currently highest bidder. Market orders are *filled instantly* (conditional on sufficient demand).
    2. **Limit order**: sell a specified amount of the good at the price that is limited by a specified minimum limit. If there are units that are demanded at a higher price, various allocation rules can be used, but usually the difference is split evenly between the buyer and the seller. Limit orders are *filled at the earliest opportunity*, and depending on availability, may only be filled partially. Therefore some limit orders may allow an *all-or-nothing option*. Furthermore, limit orders may be placed for a certain time period, after which they are automatically cancelled, or they may be placed as *good until filled or cancelled*.
- Double auctions are often used for trading in highly liquid markets with a large number of participant. For example, it is used by the New York stock exchange for intra-day trading.
- However, unless there is a high volume of continuous trading, few traders may be willing to submit market orders since the price at which objects are acquired or sold is very uncertain. Therefore will less trading, we would mostly expect to see limit orders.
- In case there is a low trading activity to run continuous markets, a restrictive form of double auction, called a **call market** may be utilized:
  - Only limit orders are allowed.
  - The orders are collected for a certain time period. Notional demand and supply schedules are constructed based on these limit orders and the market is cleared once in a while.

- Early double auction experiments: **Smith (1962, 1964)** in reaction to the pit market of **Chamberlin (1948)**
- Comparison of performance of double auction and call market:
  - **Smith (1982)**: multiple periods, fixed valuations of buyers and costs of sellers; price convergence to the competitive level generally faster and more reliable in the double auction.
  - **Friedman (1993)**: similar setup, finds little difference in convergence to the competitive price between the two trading mechanisms.
- Exercise of market power on the margin in call markets:
  - **McCabe, Rassenti and Smith (1993)**: subjects sometimes try to do that (sellers inflate their costs and buyers underreport their valuations for marginal units), but the overall level of efficiency is still high and comparable to the double auction.

FINANCIAL MARKETS AND PRICE BUBBLES

- One important setting in which the double auction is used is financial markets, particularly the secondary stock and bond markets. In recent time, the dominant line of thinking about the equilibrium price in these markets is the **efficient market hypothesis** which claims that security prices efficiently aggregate all available information and therefore truly reflect the **fundamental value** of each particular security. The argument is based on the assumption that if a market participant has convincing information that the true value of the security is likely to be either higher or lower than its current price, then this participant will engage in arbitrage that will push the price in the direction toward the fundamental value.
- This argument has been criticized heavily over time. The main theoretical line of criticism, spelled out by **Shleifer and Vishny (1997)**, goes as follows. Even if a particular market participant may see that, given his information, the security is clearly mispriced, he will not necessarily trade on this information because:
  1. The amount of trading he can do is too small relative to market turnover, so his trading actions are unlikely to affect the market price; but then it is unclear when in the future the price would converge to the fundamental value, and hence if the trader's investment horizon is short enough, it is more important to guess where the market price will be within this horizon rather than where it currently stands relative to the fundamental value.
  2. The trader may wonder whether his information may be biased given that the market at large thinks otherwise. that is, one may use the current market price to infer what the others think on average about the fundamental value of the security. If it is too different from the trader's personal assessment, it may be because his private information is biased.
- The idea also received a lot of criticism during the recent financial crisis with many people pointing out that mortgage-backed securities were clearly overpriced relative to their fundamental value.

- In real markets, however, the evaluation of the efficient market hypothesis is hampered by the fact that there is no objective way of ascertaining the fundamental value of the security. Hence any time path of prices is in principle consistent with the efficient market hypothesis because on can always interpret price swings and "corrections" as sudden changes in the fundamental value in reaction to new information.

- This controversy is what led researchers to design laboratory studies in which it is perfectly clear what the fundamental value of the security is at any point in time, and to compare this fundamental value with the actual market price in order to identify propensity of markets to create **upward or downward bubbles**.

- Impact of trading experience on the incidence of price bubbles:

  - **Smith, Suchanek and Williams (1988)** and **VanBoening, Williams and LaMaster (1993)**: sessions with experienced subjects (subjects who participated in an asset market experiment before) produced many fewer bubbles than markets with inexperienced subjects.
  - **Peterson (1993):** uses super-experienced subjects (subjects who participated in two asset market experiments before); finds that the market price tracks the fundamental value of the asset closely.
  - **Dufwenberg, Lindqvist and Moore (2005)**: if at least one third of the traders are experienced, bubbles are largely eliminated.
  - **King et al. (1993), Haruvy and Noussair (2006), Noussair and Tucker (2003); Porter and Smith (1995)**: futures markets and short sales help is reducing the incidence of bubbles

Smith, Suchanek and Williams (1988)

- A "stock" is being traded for 15 periods.

- The stock is worthless at the end of the horizon, but before that, it pays a random dividend at the end of each period. The dividend can be 0, 16, 24 or 60 with equal probability. Hence the expected dividend is 25 per period. Given the shortness of the horizon in a lab experiment, one would not expect any discounting, and hence, under risk-neutrality, the fundamental value of the asset at the beginning of period $t \in \{1, ..., 15\}$ is $25(16 - t)$.

- At the beginning of the experiment, each trader is endowed with a certain number of shares and a certain amount of cash (there are many different parametrizations; see the paper for details). Therefore one can be a net buyer or a net seller of the shares from the very beginning.

- Short sales or leveraged buying are not permitted.

- In each round, the double auction is used as a trading mechanism (like on NYSE). Each trading round lasts a maximum of 240 seconds.

- Result: in 14 out of 22 sessions, there is an upward price bubble.

- Here is a typical result for a session with inexperienced subjects and with an expected dividend of \$0.16:

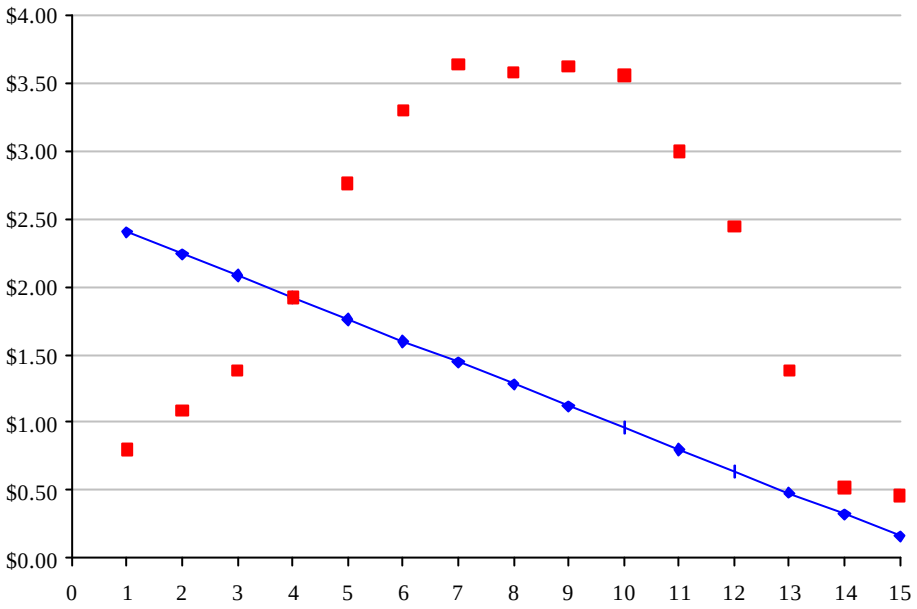

Figure 18.1. Average Transactions Prices in a 15 Round Double Auction with a Declining Fundamental Value

- The fundamental value starts at $15 \times \$0.16 = \$2.40$ in period 1 and declines linearly to \$0.16 in period 15. The market price is first below the fundamental value (downward bubble), and then above it for an extended period of time (upward bubble). At the end of the horizon, the price crashes toward the fundamental value.
- Such bubbles were observed with undergraduates, graduate students, business students, and even a group of commodity traders.
- Message of the paper: asset markets are prone to generate bubbles even if there is no uncertainty about the fundamental value (given distribution of dividends and a fixed endpoint value). Hence, the real-world markets with unknown distribution of dividends and without and endpoint may be even more prone to price bubbles.
- This result has been replicated many times ever since and seems to be fairly robust (a possible exception being when traders are educated in economics and/or finance, in which case the price path tracks the fundamental value).
- The asset market setting of this paper has ever since been challenged, though, as not being representative of the real-world asset markets. Common criticisms run along two lines:

  1. In the real world, fundamental values of assets are usually non-decreasing over time.
  2. Adjusted for purchasing power, there is no strong trend in real-world asset markets in the ratio of cash and the market value of securities. In the setting of the current paper, this ratio grows strongly over time, which may contribute to the formation of the upward price bubble in later rounds. For more work on this issue, see **Caginalp, Porter and Smith (2001)**, for example.

Bostian, Goeree and Holt (2005)

- This is one of the attempts to deal with the issue of the declining fundamental value in Smith et al. (1988).
- The idea is to have a constant fundamental value, which requires two things:
  1. a positive redemption value of the security at the end of trading
  2. discounting, i.e., a positive interest rate on cash
- In particular, the authors set the dividend process and the terminal redemption value in such a way so that the latter is the expected net present value of the security if this security is to be held until indefinite future, collecting dividends in each period.
- Also they change the market mechanism from double auction to call market.
- Baseline parametrization:
  - In each period the dividend is either \$0.40 or \$1 with equal probability. Hence the expected dividend is \$0.70.
  - The interest rate on cash is 0.1 (10%). This implies that the present value of the asset held until indefinite future is 0.7/0.1=\$7.
  - The terminal redemption value of the asset is therefore set at \$7.
  - 12 traders, each initially endowed with 6 shares and \$50 in cash.
  - 20 periods.
- Results for two typical sessions:

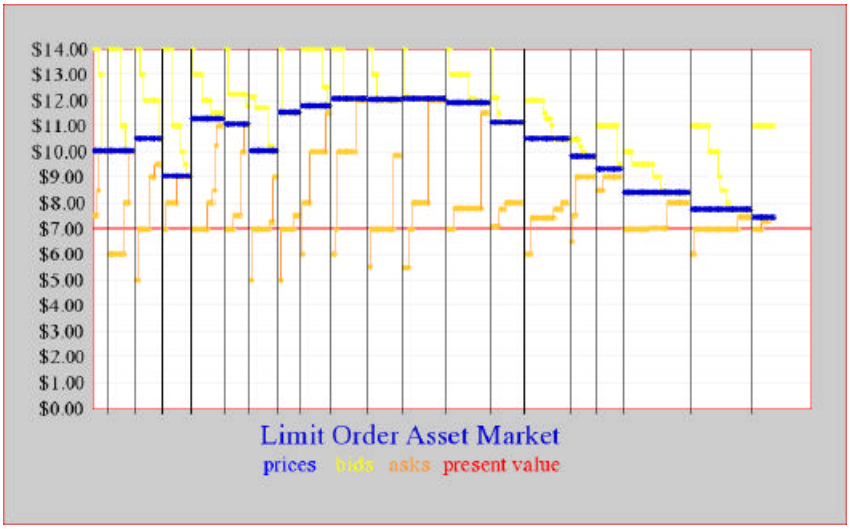

Figure 18.2. A 20 Round Call Market with a Constant Fundamental Value of \$7.00

- The authors also considered a few sessions with the horizon of 40 periods. In one of them, they recorded an extreme bubble:
- Message of the paper: the incidence of bubbles does not seem to be driven by the time-path of the fundamental value. Bubbles in experimental asset markets occur not only with declining, but also with constant fundamental values.

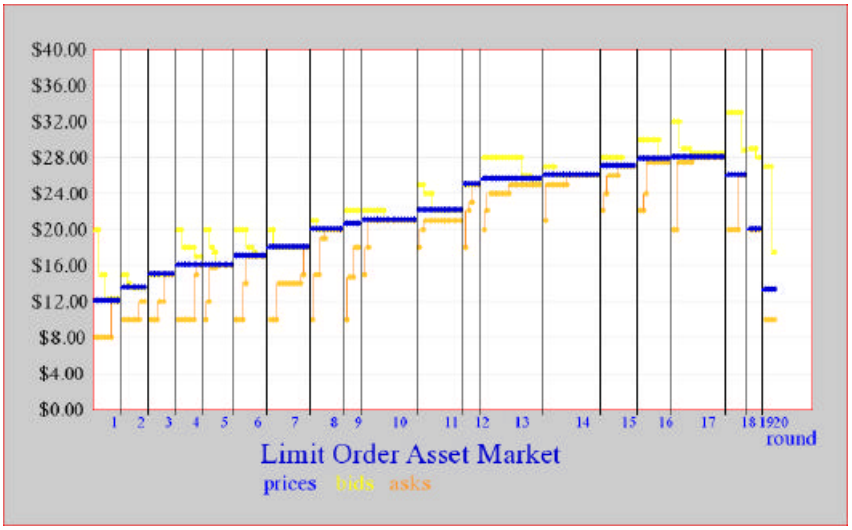

Figure 18.3. A 20 Round Call Market with a Constant Fundamental Value of $7.00

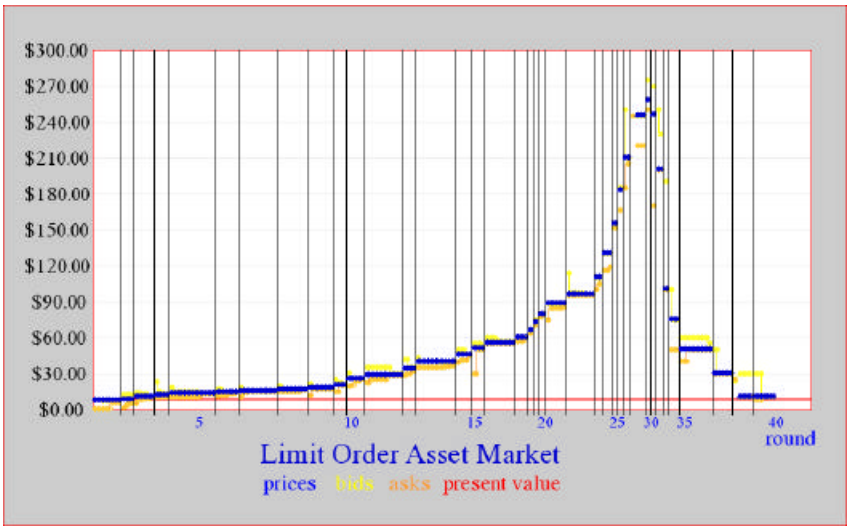

Figure 18.4. A 40 Round Call Market with a Constant Fundamental Value of $7.00

PREDICTION MARKETS

- The objective of a prediction market is to efficiently aggregate information about a likelihood of a particular event happening that is diffused in the population.
- A prediction market works as follows:

  1. The state space (space of possible relevant events) is partitioned into a set of mutually exclusive and exhaustive events.
  2. The house (the exchange organizer) creates a set of securities whose payoffs are dependent on which event is eventually realized. Each such security pays off $1 if the corresponding event is realized and $0 if some other event is realized. Such security is called an **Arrow security**.
  3. The house stands ready to exchange $1 for a market portfolio (a portfolio consisting on one of each of Arrow securities) in either direction in any quantity at

any point in time. (Note: the house may charge a lump-sum participation fee or a small surcharge for cash inflows in order to generate its own revenues.)

4. There is a double auction market for each of the securities. Clearing may be continuous or it may be done at discrete time points.

- **Idea**: based on the efficient market hypothesis, market prices of the individual Arrow securities (which will always be between 0 and 1, inclusive) reflect average market beliefs about various events happening.
- **Manski (2006)** doubts whether equilibrium prices will indeed reflect average market beliefs, but **Wolfers and Zitzewitz (2006)** refutes the argument and proposes a model that shows that equilibrium prices will indeed be good approximations of average market beliefs.
- The most famous example from the field: **Iowa Electronic Market** as a predictor of outcomes of US presidential elections. Furthermore, **Wolfers (2004)** and **Berg, Nelson and Rietz (2008)** document that predictions from this market outperform poll-based predictions three quarters of the time.
- Other uses: to forecast likely probability of political events or range for sales growth or similar figures within companies. Idea: using anonymous market in which disclosing information via trading pays off, individuals may be tempted to reveal information that they otherwise would not reveal (due to the incentive structure).
- Various examples of use (taken from **Wolfers, 2004**):
  - The Department of Defense in the US considered running a Terrorism Prediction Market (the idea was abandoned soon, though).
  - Goldman Sachs and Deutsche Bank started prediction markets on employment, retail sales, business confidence or inflation.
  - tradesports.com offered contract for the event that Saddam Hussein would be ousted from power by the end of June 2003.
  - Hollywood Stock Exchange runs prediction markets on Oscar winners.
  - A prediction market at HP produced more accurate estimates of printer sales than internal firm forecasts.
  - A prediction market at Siemens correctly predicted that a deadline would not be met on a particular project, even if internal planning suggested that it could.
  - TradeSports and World Sports Exchange frequently run prediction markets for regional and local elections.

Figure 2
The Saddam Security

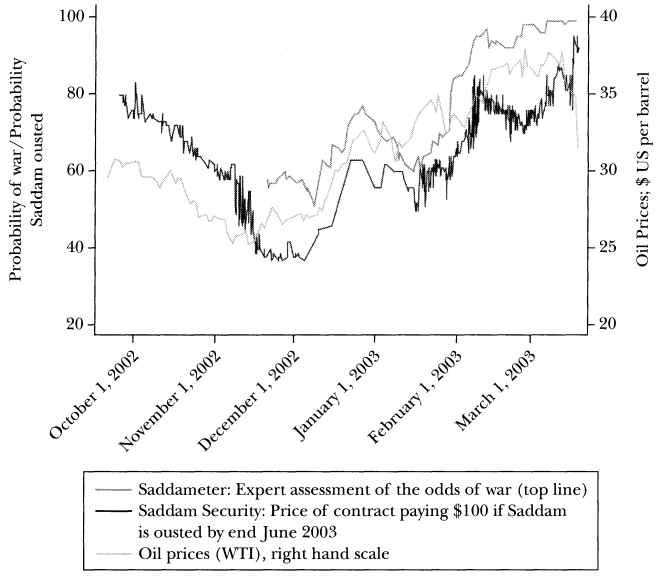

Sources: Trade-by-trade Saddam Security data provided by Tradesports.com; Saddameter from Will Saletan's daily column in Slate.com.

Figure 3
Predicting Movie Success

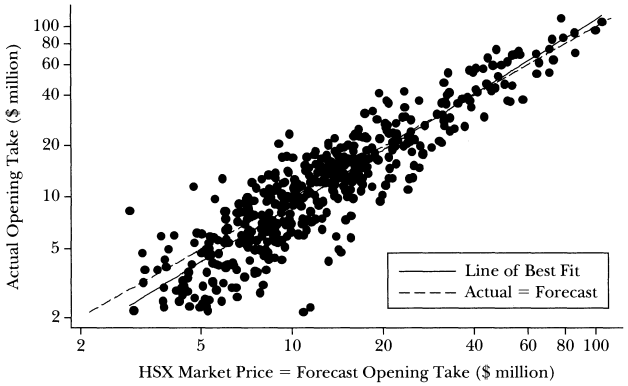

Source: Data from 489 movies, 2000–2003 (⟨http://www.hsx.com⟩).

- Example from home environment: **Katarina Kalovcova**'s work on prediction markets for:

  1. the number of applications to the CERGE-EI Ph.D. program in the spring of 2009 (state space partitioned as 81-90, 91-100, ..., 171-180, tournament with prizes for the three best performers, no buy-in)
  2. the number of applications to the CERGE-EI Ph.D. program in the spring of 2009 that come from Czech or Slovak Republics (state space partitioned as 4-5, 6-7, ...25-26, 300Kc buy in, payoffs equal to actual performance)

  – Call market running for several weeks, clearing once or twice a day, respectively.
  – Information on the number of received applications up until then was available on a daily basis from the Student Affairs Office.
  – However, participants may have had private information about friends or friends of friends who were about to submit their application and they hence may have

Figure 4
2003 California Gubernatorial Election

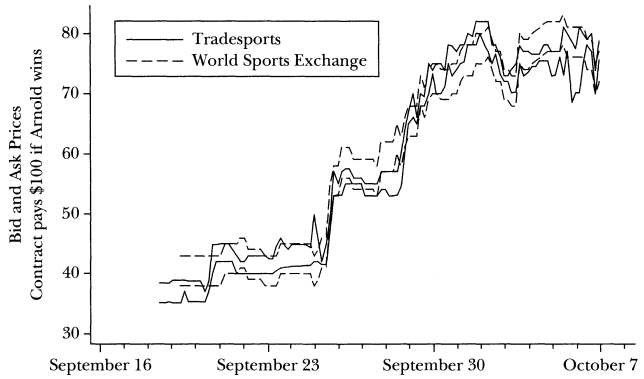

Source: Prices collected electronically every four hours by David Pennock.

commanded an informational advantage, especially toward the end of the trading horizon.

- Results: the market for the overall number of applications, due to the tournament structure of payoffs, invited a lot of risk-taking, narrow portfolio focus and collusion. Although the market tracked the developments in the number of received applications fairly well, there was a lot of noise.
- The market for the number of applications from Czech and Slovak Republics did much better given its payoff structure.

- What we have discussed up until now are referred to as trades in **winner-take-all contracts**. Another possibility for the design of prediction markets is to trade **index contracts**. In such contract, each security pays based on a percentage that a particular event/candidate receives. Payoffs of these securities therefore vary continuously with vote proportions, for example, unlike payoffs of winner-take-all contracts.