Chapter 5

Economic Models

CONTENTS

Several varieties of social computing, especially crowdsourcing, rely upon motivating people to participate in such a manner as to produce desired value. This value may be reflected in the amount of personal effort they expend, the additional resources they bring to bear upon the problem, and the quality and truthfulness of the information they provide.

The economic abstractions provide a well-established approach to capture interactions of autonomous parties who can be modeled as being self-interested and strategic in that they can choose their actions, that is, exercise their autonomy, in a manner so as to improve their payoff. They can be particularly effective where we are given multiple parties with divergent interests albeit in settings where we can control their interactions sufficiently to impose some integrity constraints on the interactions (for example, cases of business competition but not necessarily battles unto death) and where we provide appropriate incentives (taken here as including disincentives) to the parties.

5.1 Mechanisms and Mechanism Design

In the sense of economics, a *mechanism* is a set of rules of encounter. That is, a mechanism defines the economic environment under which the participants operate. Specifically, a mechanism describes how the interactions in a society are regulated. These includes how any disputes may be adjudicated and how participants are encouraged to undertake prosocial behaviors.

A large variety of mechanisms can be defined and have been defined. These include *(a)* Honor systems, which presume that each participant will act honorably according to social customs; *(b)* Honor systems with social censure (as a penalty): these are more typical since the existence of social mores on honor suggests the possibility for censure for violating them; *(c)* Spot checking as typically used for paying taxes, which are voluntary or self-declared actions, but with selective audits and severe penalties for violators; and *(d)* Auctions and markets, which focus on economic incentives as a means to regulate the relevant behavior of participants.

In general, mechanisms describe the outcomes that affect participants under specified circumstances. In some cases, the outcomes could be directly stated as rewards and penalties but in other cases not. For example, a market may specify what price a participant must pay for an item—based usually on what this and other participants are willing to pay or accept. If the participant values that item more than this price, the outcome would be a reward for that participant; if the participant values that item less than the price, it would be a penalty.

It is instructive to compare mechanisms such as the above with respect to properties such as whether they produce high social welfare and how robust they are against violators. For example, we can see that the plain honor system has extremely low overhead since each participant presumably does the right thing but it is also highly fragile in that it is defenseless against violators—dishonorable people who will cheat with impunity. Adding in social censure increases the potential for compliance but also increases the overhead. People have to keep tabs on each other and be willing to criticize them to their faces or badmouth them behind their backs. Since there is no formal process for catching violators, the sanctions may be wrongly applied. In such a case, the mechanism can lead to a miscarriage of justice and appear quite unfair and unethical to those who are inadvertently victimized. Spot checking requires a more elaborate process to determine potential violators; the penalties must be high enough so that there is an effective disincentive for violation. In that case, the process to verify that there is a violation must be sufficiently rigorous to ensure that the severe penalties are not wrongly applied. The economic mechanisms provide incentives (and disincentives) to the participants—in general, encouraging them to exert the necessary effort and reveal the desired information. They often rely upon capturing from the participants as a basis for computing the incentives. They rely upon a robust architecture, specifically, a market (a server in computational terms) that gathers information, takes decisions regarding the outcomes and incentives, and enforces those decisions. The enforcement is crucial—an earned reward must be provided and an earned penalty must be applied. Ensuring the credibility of the market is essential to its effectiveness.

For this reason, when we approach social computing from the economic standpoint, the problem becomes one of *mechanism design*. That is, how can we specify a mechanism that would provide autonomous participants the appropriate incentives so they would willingly act in a manner that promotes the goals of the designer. These design goals can include various system-level properties, such as that participants interact productively and fairly. For example, if we are interested in obtaining a true estimate of the weight of an ox, we should offer a payoff to the participants whereby they would stand to gain from producing a true estimate.

5.1.1 A Puzzle from the Talmud

As an illustration, let us consider an ancient puzzle derived from a story in the Talmud. But, first, let us consider a simpler setting. Suppose you are given two men each with one horse and wish to determine which of the two horses would run a mile faster than the other. Assume here that the men are autonomous and self-interested, whereas their horses are merely their instruments who would do their bidding.

To make the above determination, we can create a short race (for a mile precisely) and offer a reward to the owner of the horse that wins the race. Since the owner would win a reward by establishing that his horse is faster than his competitor's horse, the owner would act in a manner that would demonstrate his horse's capabilities in the best possible light. For example, each owner would race his horse at the fastest speed that the horse is capable of in a one-mile race. (The owner may try to hire the best jockey available for the race but let's ignore such complexities.) In this example, the mechanism consists of these parts: *(a)* a setting: the race between the horses; *(b)* an action by each party: how well to ride in the race; and *(c)* a reward: the winner of the race wins.

Along with a mechanism we have strategies by which the participants make their decisions (such as about the actions they perform). In the above case, the owners can have various strategies. Relevant to the proposed competition, two of the possible strategies are to *(a)* try one's best to win; and *(b)* not try one's best to win. The mechanism is such that it is in each owner's interest to try to win (so as to obtain the reward). This strategy dominates the strategy that involves not bothering with winning. Arguably the strategy to try hard to win dominates any other strategy of the owner. Please note that in this case we are simplifying the setting so that each owner must participate or that the prospective reward is sufficiently high from the standpoint of each of the owners that both of the owners will participate. In more realistic settings such might not be the case. A rational party may decide not to participate in the mechanism because the party's expectations regarding the odds of obtaining a large enough reward are low. Section **??** returns to this point.

The foregoing mechanism makes good sense for the above problem. Now let's consider a small variation of the above—this variation is from the Talmud. Here we seek to determine not the (owner of the) faster horse but the owner of the slower horse (say, over a one-mile run). Let's consider what would happen if we naïvely modified the previous mechanism to tackle the new case. We would conduct a mile-long race between the two horses but would offer a reward to the owner of the horse that finished second. Such a mechanism is obviously flawed. Since winning the race is evidence that one's horse is the faster of the two, each owner would try his or her hardest to lose the race. In particular, the owner would not instruct its horse to exert. In the limit, if each owner tries to go slower than the other owner, the race wouldn't terminate—or, rather, won't get off the ground.

It's worth pausing a little to think about how you would design a mechanism so as to tackle this puzzle.

The solution that the Talmud offers builds on the idea that we must align the incentives of the participants—the owners of horses—with the design goal of rewarding the owner of the slower horse. Since each owner wants to win the competition and can win the competition when his opponent's horse is shown to be the faster of the two, a simple solution is to put each participant in command of the other's horse. That is, the owners ride each other's horses in the race. Now each owner rides the horse he is on as hard as he can to try to win the race because winning the race on the other owner's horse is the only way to win the competition.

5.1.2 A Pathological Mechanism

An interesting—and pathological—mechanism was proposed by Martin Shubik and colleagues, reputedly including the celebrated Nobel laureate John Nash [Shubik 1971]. Shubik's Dollar Auction involves the sale of a dollar bill: the bill is of no special significance—that is, not a historical bill or whatever. Bidding begins at 10¢ and continues in increments of 10¢. The auction ends when no one comes forward with a higher bid than the current highest bid. Like most auctions, the highest bidder wins the auction and gains the dollar bill. However, unlike other auctions, here the highest and the second highest bidder each pays the respective amounts that they bid.

For simplicity, consider a setting where we have two prospective bidders, Alice and Bob. Alice notices that a bid of 10¢ would net her a profit of 90¢ if she wins, so it is rational for her to bid 10¢. Bob notices that a bid of 20¢ would net him a profit of 80¢ if he wins, so it is rational for him to bid 20¢ and take the lead in the auction. Similarly, each party perceives a gain (although diminishing in each round) and keeps bidding. Let's consider the case where Alice's bid is 90¢ and Bob's bid is 80¢. If Bob quits, he has a guaranteed loss of 80¢. By upping his bid to 100¢, he will only put in an additional 20¢ but will gain a dollar. That is, he would still net 80¢ by bidding. But once Bob takes the lead, Alice is exactly in the same position. In other words, each party has a myopically rational strategy of continuing to bid ever-larger amounts for a dollar—even more than a dollar.

Shubik has described situations [Poundstone 1992, pp. 262–265] wherein he would use this auction as a parlor game for guests at dinner parties. Shubik observed that his guests often bid well above a dollar because they would get into a bidding war with someone else. They would stop only in frustration and with their feelings hurt.

For our present purposes, Shubik's auction is an example of a mechanism that drives rational parties into irrational decision making. In essence, it creates the possibility of retaliating against someone. With only a small amount of investment each time—the additional 20¢ you need to risk to go from your previous bid to your next bid

5.2 Economic Abstractions

The economic abstractions are based on the idea of the parties involved being rational in the sense of having their motivations characterized by preferences that they seek to fulfill or, equivalently, by a numeric utility that they seek to maximize. The economic abstractions are powerful in that they support achieving optimal resource allocations in a way that addresses the private goals or preferences of the concerned parties.

However, these abstractions are incomplete by themselves in that they presume *(a)* sufficient stability in the environment that economic transactions can proceed; and *(b)* that the motivations of real-life social entities (such as people, enterprises, or organizations) can be characterized in economic terms.

These abstractions provide a basis for achieving contracting behaviors, specifically, in helping the parties to arrive at contracts acceptable to both of them through some variant of negotiation procedure. The economic abstractions additionally characterize how each party ought to behave given its preferences and the applicable mechanisms.

To realize the economic abstractions involves some crucial assumptions. In particular, the traditional economic theories assume that the parties are perfect reasoners and can think up the best strategies for themselves given the mechanisms that apply.

Let's consider the problem of sharing resources between two or more people. Suppose the resource is a cake: a heterogeneous dessert with icing and filling. For simplicity, let's consider exactly two people. The relative sizes of two pieces and other attributes (e.g., the amount of icing on a piece) can vary. What is the correct way to split such a resource? Clearly the preferences can vary. For example, Alice may prefer a smaller piece as long it has more icing than the other. There is no method, central or distributed, that would guarantee that both Alice and Bob leave happy.

The traditional children's method of "I will cut, you choose"—in which one child cuts the cake into two pieces and the other child chooses one of the pieces—is remarkably good for two-player settings. However, that method would not ensure that the first child found the best cut possible that would have produced pieces that each of the children would have liked.

5.2.1 Valuations

An elementary question that the economic abstractions help understand is how trade can possibly work. Regardless of the means of trade—whether bartering goods (and services for that matter) or using money—why would rational players voluntarily participate? How can two trading parties both gain from a trade? And, if they don't, why would the losing party autonomously decide to participate? Trade, of course, occurs. So the problem lies in one of our tacit assumptions, namely, that if one trading party gains from a trade the other trading party must suffer. This assumption is fallacious, The valuations of the goods involved are not centrally determined. Thus, it is possible that two trading parties may value the goods being traded differently. Their difference in valuation reflects their autonomy as well as the difference in their respective circumstances. It also provides for circumstances where both trading parties may gain. As an example, consider trading a bottle of drinking water for USD 1. The vendor may have numerous bottles of water and may have no interest in consuming them himself but an active interest in obtaining money. The vendor may value a bottle of water as being below USD 1, let us say at USD $1-\delta$, for some $0<\delta<1$. The buyer may have money and a need for bottle of water and may value a bottle of water as greater than USD 1, let us say at USD $1+\varepsilon$, for some $0<\varepsilon<1$. Now, if the trade takes place with the vendor providing a bottle of water to the buyer and the buyer providing USD 1 to the vendor. Both parties gain from this trade at the price of USD 1: the vendor gains δ and the buyer gains ε.

In a nutshell, it is the independent valuations of the goods being traded—or, rather, differences in such valuations—that enable trade. Of course, it is possible for the valuations to be such that both parties would not gain from a trade. In that case, we consider the current allocation of resources as being optimal.

How do players place values of goods? In general, coming up with valuations is a complex process that incorporates considerations of one's genuine needs, psychological makeup, and social influences. A common simplifying assumption is that valuations are *independent and private*, meaning that how much value one social entity places in some good (or, generally, outcome) depends solely upon that social entity and is unaffected by others. An example would be something like cake or ice cream that a person would consume or use and whose consumption or use of it is not influenced by others. The assumption is that the pleasure someone derives from eating a piece of cake would be unrelated to the pleasure someone else derives from eating that piece of cake. Indeed, it is possible that the cake produces pleasure for one person but is unpleasant for another (for instance, if the second person has an allergy to some ingredient of the piece of cake).

The existence of independent, private valuations is an assumption that can be questioned on psychological grounds. All too often, our desires—the valuations we assign goods—are shaped by the desires exhibited by those around us. A little girl may want a vanilla icecream if she sees her older brother have a vanilla icecream. A boy may want a toy he sees other children playing with.

For some goods, however, it is appropriate even on grounds of economic rationality to assume that the valuations are *common*, meaning that the valuation one party assigns to a good is entirely based on the valuations of others. For example, we might expect that no one inherently values treasure bills or currency (whether local or foreign): for such a good, its valuation equals what other people think it is worth.

In general, however, valuations are neither entirely independent nor entirely common but are *correlated*, meaning that they are a combination of the above. Real-estate property such as house and even an automobile have this trait. Part of the valuation for a house derives from the pleasure a home owner would derive from living in it, for example, enjoying a view of the mountains. Part of the valuation would derive from its potential resale value, which depends upon how other people would value that house.

A important tenet of valuations that it does not matter whether a good is currently possessed by someone (possession, of course, determines whether a trade can take place).

5.2.2 Allocations

We can think of goods and money uniformly as resources. Money is a particular kind of resource in that its valuation we can assume is explicit. Economists refer to such a resource—in terms of which other resources are valued—as a *numéraire*.

An *allocation* is mapping of each resource to exactly one social entity. That is, we can think of an allocation as a vector with one element for each participant, describing the resources that that participant owns. Allocations are important because many social computing questions can naturally be formulated in terms of resource allocation.

What is an optimal or correct resource allocation? Since we are concerned with two or more social entities and finite resources, it is clear that their interests would not be aligned. Whatever resource is possessed by one entity is not possessed by another. For this reason, we cannot arbitrarily decide who should get what resource.

One of the most established standards for the correctness or optimality of a resource allocation is Pareto optimality named after the economist and social theorist, Vilfredo Pareto. Pareto's idea was to consider allocations in terms of the valuations of the various resources by the various social entities. Thus we can evaluate resource allocations not in terms of the raw resources but in terms of the valuations of those resources. We can assume the parties involved are self-interested but not malicious. That is, being rational, each party would like to satisfy more of its preferences—in simple terms, each party would prefer a greater valuation over a smaller valuation. However, being cooperative, each party would not object to others gaining if the gains of others did not come at any loss to its valuation.

With this motivation, Pareto defined one allocation as being superior to another allocation if the first allocation is better for at least one party and no worse for any party. Now, an allocation is optimal provided there is no allocation that is superior to it. Given a Pareto optimal allocation, we can improve the circumstances for one party only by causing a loss to another party. In other words, the only ways to modify an optimal allocation would involve the use of coercion on one or more of the parties. In contrast, for a nonoptimal allocation and only for a nonoptimal allocation, we can improve the circumstances for one party without causing a loss to anyone.

Let's consider that Alice has a good g presently, which she values as \$1. Bob values the same good g at \$3. Suppose Bob buys g from Alice for \$2. In the resulting allocation, Alice has \$2 more than she did earlier but no g, whereas Bob has \$2 less than he did earlier but has acquired g. Clearly both have benefited from the trade, netting a gain of \$1 each. In fact, they would both benefit for any price strictly between \$1 and \$3. This means that the initial allocation was not Pareto optimal. In this case, the resulting allocation arose from a trade that would benefit both Alice and Bob. Now consider an alternative history in which Bob had stolen g from Alice (and thus paid her nothing for it). In the resulting allocation, Alice has no more money than she did earlier but has lost g, whereas Bob has the same amount of money as he did earlier but has acquired g. Clearly, Alice would not have voluntarily gone from the initial allocation to this one, since it would involve her giving up g for free. Yet, this allocation too is Pareto optimal. It is optimal because there is no way to improve the allocation any better for both of them.

How do we find Pareto optimal allocations in general? Under the assumption of no central control (that is, coercion), the only approach we have is for the parties themselves to reallocate their resources. They would do so by trading resources with one another. In this terminology, a nonoptimal allocation is one in which gains can be made through trade.

As described above, trade is generally possible when the (independent, private) valuations of the parties differ and those in possession of a resource value it less (with respect to another resource) than those who are not in possession of that resource. For example, suppose Alice has ice cream and Bob has cake. Suppose also that Alice prefers cake to ice cream and Bob prefers ice cream to cake. In this case, our initial resource allocation is not Pareto optimal because if Alice and Bob were to trade ice cream for cake, both would be better off (as judged by themselves individually). In addition, once they do trade, no further gains from trade would be possible. Hence the resulting allocation would be Pareto optimal.

In general, the purpose of markets is to help find Pareto optimal allocations. Assume for the moment that each participant in a market reveals its true preferences regarding the various resources under consideration. In a market, a participant's preferences are phrased in terms of the prices that the participant associates with each resource. If a party that doesn't currently possess a resource places a greater price on it than one that possesses it, the two of them can trade; each of the parties would gain from the trade provided the price of the trade lies between their respective prices.

A market helps achieve Pareto optimal by making sure that the party that preference a resource the most ends up possessing that resource. Given a resource, the party with the lowest valuation that possesses the resource trades with the party with the highest valuation that does not possess the resource. This holds for each resource in the market.

Whether a participant happens to possess the resource in question is incidental to the idea of preferences. Buying and selling are symmetric—the only difference being the current possession of the resources.

Now the success of a market in achieving a Pareto optimal allocation depends upon the participants having revealed their true preferences. In general, the parties need not reveal their true valuations. However, we can consider mechanisms that induce party to reveal its valuations truthfully.

5.3 Markets

The most established economic mechanisms are based on the notions of markets. A market is a setting (physical or virtual place) where participants can offer goods or money to one another and which provides a mechanism through which desired contracts can be created and executed.

In our common experience, markets can be of different types. The following are some well-known dichotomies. First, a market may be public or private, that is, part of a restricted exchange. Second, a market may restrict the kinds of goods traded. An *endogenous* market deals with goods that are defined within the market. An example is of a typically stock exchange such as NASDAQ. An *exogenous* market, such as eBay, provides a setting where not only virtual but also physical goods can be traded outside the scope of the market. In general, a market should provide a basis for the ensuring the integrity of the interactions of the participants in it. A specific example would be some form of nonrepudiation. Third, markets can differ in the variability of the goods traded in them and knowledge of the prices of those goods. For example, a stock market deals with securities but the securities are mostly easily judged to be of the same type. The real-estate market deals with goods each of which is fundamentally unique. For example, two shares of IBM are mutually indistinguishable, whereas any two three-bedroom houses are distinct (with different locations, different neighbors, different sunlight in the kitchen, different school districts, and so on).

An interesting property of markets that makes them magical to some is that under certain assumptions, markets ensure a resource allocation where all the goods that can be sold (for the market price) are sold and all the goods that can be bought (for the market price) are bought.

5.3.1 Pricing

The general motivation is to allocate resources to those who value them the most. Pricing is a way to help accomplish this.

Fixed pricing refers to prices being fixed or, rather, slowly changing. Prices could be set based on any appropriate criteria, such as *(a) flexibility*, for example, restricting rerouting or refundability in air travel; *(b) urgency*, for example, a convenience store can charge more than a warehouse; *(c) customer preferences*, such as using coupons to attract price-sensitive customers whereas others pay full price; *(d) demographics*; *(e) artificial*, for example, the so-called Paris Metro or Delhi "Deluxe" bus pricing where there is no fundamental gain in quality other than that the increased pricing itself helps reduce demand and thereby reduce congestion; and *(f) predicted demand*, such as prices being different at peak versus off-peak times of the day or week, as on the New York subway, or in phone and electric power rates. As an interesting bit of trivia is that the idea of time-dependent pricing was introduced for the New York metro at the suggestion of William Vickrey, whose eponymous auction we study elsewhere in this chapter.

Dynamic pricing is when the prices rapidly changing, based on current demand and supply. Participants seek to maximize individual utility; those who value something more will pay more for it. Fixed pricing leaves some revenue that other parties exploit (e.g., in secondary markets such as black markets as in football and music concert ticket scalping). Musicians in particular would like to fight back against [*Iron Maiden Take up Fight against Online Ticket Scalpers* 2015]. They face a dilemma in that they would like to price their tickets low enough that all their fans can afford to attend their concerts but since some fans are willing to pay a lot more than others there is an inevitable opportunity for scalping [Davidson 2013].

5.3.2 Centrality of Prices

A price is a scalar, such as an integer or a real number. The important point is that prices can easily be compared: given two prices, we know which one is bigger or whether they are equal.

A key assumption is that the computational state of a market is described completely by current prices for the various goods being traded in that market. That is, the trading parties do not communicate with each other in the classical public markets, where all information is transferred through prices.

Communications occur between each participant and the market as a computational and a social (that is, autonomous) entity in its own right, and these communications consist only of identifying the goods to be traded as well as the prices that they charge. Therefore, in these idealized market settings, the participants reason about each other and choose their strategies entirely in terms of prices being bid.

But of course in ordinary settings the absence of communication is not a viable assumption. For example, in real-estate markets, buyers and sellers can interact in person or through their respective real-estate agents. In other settings, such as in the canonical fish markets, haggling is common. Buyers and sellers can try to influence each other about the prospects for prices because of expected events, such as the weather. Buyers and sellers may respectively disparage or praise the products being bought and sold. Trash talk is a time-honored tradition not only in sporting contests but also in markets.

5.3.3 Equilibrium

A market is in *equilibrium* when supply equals demand. At equilibrium, the market has computed the allocation of resources. Equilibrium carries the connotation of having computed a market price. This is the most direct connection between a market and social computing. The intuition of balance that a market equilibrium connotes is exactly the same as what Galton advocated in his *Vox Populi*. The number of items people are willing to sell at prices low enough to be sold at the equilibrium price equals the number of items people are willing to buy at prices high enough to be bought at the equilibrium price.

One of the early and significant results about markets is that they compute a simultaneous equilibrium (of supply and demand) across all resources. That is, the prices of the resources change until equilibrium is achieved and the market *clears*. In the late 1800s, the French economist Léon Walras, teacher of our friend Vilfredo Pareto, analyzed the flow of goods in and out of Paris and was intrigued by how everything appeared to balance out every day.

This study led Walras to formulate the notion of how a market—or, rather a set of markets, one for each good—clears when the prices for the various goods move up and down dynamically until supply equals demand for each of them. This occurs just as easily as when the goods are apparently unrelated to one another as when the goods are closely related and intersubstitutable with one another. For example, we might imagine that customers who buy brown eggs one day would be less interested in white eggs the same day: thus the supply of brown eggs can affect the demand for white eggs. But the prices of the two types of eggs would move up or down depending upon which eggs are preferred by more people on a given day.

5.3.4 Computational Aspects

In computational terms, a market takes requests (in particular, buy and sell bids) from participants, enforcing rules such as bid increments and deadlines. It provides part or all of the market state (current bids for what by whom) to the participants. Lastly, a market determines the outcomes, that is, the reallocation if any of the resources based on the preferences expressed by the participants.

We can think of an auction as a computational mechanism to manage supply and demand. An auction works via a common unit of value, which we can think of as money. Several varieties of auctions are known, including those that are silent (auctioneer names a price; bids are silent) versus (bids name prices; auctioneer listens); whether they disclose or hide the identify of the participants and their bids. More interestingly, auctions can be *combinatorial* meaning that they: involve bundles of goods.

A market inherently involves matching buyers and sellers. Buy and sell bids can be matched in various ways, which support different properties.

One useful distinction is about the nature of prices. Prices can be *uniform*, meaning that multiple units, if simultaneously matched, are traded at the same price. That is, if multiple buyers and sellers match, each respectively pays and obtains the same price. An alternative is that the prices are *discriminatory*, meaning that the trading price for each pair of bidders is potentially different. For example, we might determine the price for each matching pair as the sell bid (in that pair). Or, we might set the price for a matching pair as the bid that was chronologically the first or chronologically the last.

Clearly, uniform prices are simpler to describe. This simplicity matters not only for ease of implementation but in providing a clearer description of the mechanism to the participants. In effect, under uniform pricing, the participants together compute what the market knows and thereby learn from each other. They need to spend less effort worrying what clever strategy to apply so that they can obtain the best of the multiple prices that will result from the current episode. They would suffer less from angst about whether they made the right choice in bidding the amount they did at the time they did.

5.3.5 Auctions

Let's briefly consider some well-known examples of auctions. In an English Auction, buyers bid for an item. The essential features are that *(a)* only increasing bids are accepted; *(b)* the highest bidder wins the auction; and *(c)* the winner pays the price he or she bid to acquire the item being auctioned. Some variations include a specification of the *(a)* minimum bid increment; *(b)* a reserve price, which the highest bid must exceed, else there is no sale; and *(c)* a time bound.

In contrast, in a Dutch Auction, originally invented for auctions of tulips, there is a price "clock" or counter that starts high and winds down. A bidder may stop the price clock. The first bidder who does so wins and pays the price on the clock when it is stopped. Here too, the highest bidder wins and pays the price he or she bid.

The Fish Market Auction is based on the fish market held in the village of Blanes near Barcelona in Spain. Buckets of fish are sold off in this auction. In this auction, a small crowd gathers around an auctioneer, who calls out prices for a bucket of fish. When two or more buyers bid (by calling out "Si!"), the auctioneer calls out an increased price. If no buyer responds positively, the auctioneer calls ut a decreased price. The process continues until exactly one buyer is found.

The above auctions let bidders be aware of other bidders' bids. In a Sealed Bid First-Price Auction, each bidder submits a bid without knowing what other parties are bidding. Additionally, there is only one opportunity to bid. There is therefore a deadline. The auctioneer gathers all the bids once the deadline is past. The highest bidder wins and pays the price he or she bid. Such auctions, sometimes called *tenders*, are commonly used by governments and large companies to award certain large contracts—in the case of a contract, the lowest price quote wins.

5.3.6 Winner's Curse

These settings describe a form of the winner's curse [Milgrom 1989]. Suppose a set of mutually independent, though well-informed, individuals estimate the value of some relevant good or service and bid accordingly. For example, these people could be considering the gain to be expected from a building contract. The conditions of Condorcet's Jury Theorem, discussed in Section **??**, apply. Therefore, the best estimate of the value of the relevant good (here, contract) is the mean of the estimates from the various people.

Yet, the winner of the contract would be the individual who placed the highest bid. In essence, this mechanism necessarily selects a winner who is over paying relative to the population.

A second form of the winner's curse arises from the setting of any auctions such as as the English or the Dutch. The winner of the auction would pay a price for it, say, the price is $x. Suppose the winner turns around and attempts to sell the item. Quite obviously, once the winner makes the purchase, the available buyers would include only those who placed bids smaller than the winner. Thus, if the winner immediately try to resell the item, he would find that it sells for a price lower than what he paid for it. Even if the item is inherently valuable to the winner, he could have obtained it for less than he paid.

This situation additionally indicates the possibility that the winner could try to lower his potential buyer's remorse in the purchase by trying to estimate what the next bid would be. For example, if the winner knew the next bid were $5, he would have bid not $7 but $5.01—in order to ensure that he obtained the item for the lowest possible price.

5.3.7 Vickrey Auction

The Vickrey auction is a remarkably simple yet profound idea. It was popularized and analyzed by the economist William Vickrey in 1961 and has had a deep impact not only on economics but also on social computing. Its impact on research was a significant factor in Vickrey winning the Nobel prize in October 1996, unfortunately for him the same week that he died.

The basic idea of the Vickrey auction lies in the following steps pertaining to the sale of one item.

- Each bidder submits a sealed (confidential) bid.
- The highest bidder wins.
- The winner pays the *second* highest price that was bid.

The main difference from a conventional auction, thus, is that the Vickrey auction uses the second highest price. Interestingly, this small shift has a major consequence.

5.3.8 Equilibrium Prices

Suppose we are given a market in which there are M sell and N buy sealed bids, each for a single unit of some item. What is the market price for the item in this case? This situation corresponds to a crowdsourcing scenario where we seek to determine the price for the item based on the bids placed by the participants.

As we discussed above, the fact that some bids are for selling and some for buying the item is not relevant in regards to determining the valuation of the item by the various participants. However, the difference between buying and selling is significant in determining where the items will flow when the transactions conclude.

Seller's reserve price is the sole sell bid. We can assume the minimum value, if no explicit reserve price. That is, a seller who expresses no reserve price is willing to accept any payment.

Importantly, we can determine two significant discrete prices (as well as other prices in between them) based on the buy and sell bids received by the market. Specifically, these significant prices are the *(a)* Mth highest of the entire set of bids: this is called the Mth price and generalizes the first-price auction, which corresponds to $M = 1$; and *(b)* (M+1)st highest of the entire set of bids: this is called the (M+1)st price and generalizes the Vickrey or second-price auction, which corresponds to $M = 1$.

We interpret a sell bid of \$x as a seller's willingness to sell for a price of \$x or higher and we interpret a buy bid of \$x as a buyer's willingness to buy for a price of \$x or lower.

The above prices are important in a particular way. What might make them special? To illustrate the importance of these prices, consider a simple case where all the bid prices are distinct. Suppose there are three sell bids (\$2, \$4, \$6) and four buy bids (\$1, \$3, \$5, \$7). Figure 5.1 describes this case.

Here $M = 3$. The Mth highest price among these bids is \$5 and the (M+1)st highest price is \$4. Let us consider prices from \$0 up to \$8. At \$0, the four buyers

Figure 5.1: Illustration of Mth and (M+1)st prices. Here, $M = 3$.

would be willing to buy but no seller would be willing to sell. Same situation at $1. At $2 and $3, one seller and three buyers are ready. At $4, we have two buyers and two sellers ready to transact; the same situation arises at $5. At $6 and $7, three sellers are ready to transact but only one buyer.

Recall that an equilibrium price, also known as *market price*, is one at which supply equals demand. Thus the magic of the Mth and (M+1)st prices is perhaps obvious. They delimit the range of prices at which supply and demand are balanced. That is, the range of prices is the equilibrium price range. Go above the Mth price and some prospective buyers bow out; go below the (M+1)st price and some prospective sellers fall out.

Notice that both Mth and (M+1)st price auctions are examples of uniform pricing.

5.4 Incentive Compatibility

One of the motivations of social choice theory is that it seeks to achieve democratic decision making, that is, to come up with allocations of resources in society in way that avoids having to coerce anyone: that doesn't mean everyone will get what they most want but they would, in principle, as good a shot at getting it as anyone else. In this way of thinking, the underlying decision mechanism would be something like voting in elections or participating in a market.

A common feature of these mechanisms is that they, first, provide an opportunity to independent players to express their preferences regarding the relevant social outcomes (where to build a bridge; whom to install as governor; how to determine who obtains a resource such as land or food) and, second, provide a means through which the relevant social outcome, such as a resource allocation, is determined. The determination procedure is necessarily a function of the preferences that the various parties express.

Now if it turned out the participants could state their preferences falsely—one presumes to gain an advantage over others—that would call into question the entire program of democratic decision making. For example, suppose during a water shortage a township allocated water proportional to whatever a household reported as its need. If we did so, households would gain by exaggerating their need. If one household were to exaggerate its need, it would increase its allocation at the cost of others. Thus the other households would have an increased motivation to exaggerate. There is no natural limit to such exaggeration by households. Therefore, the whole operation would collapse.

If instead we could find a way wherein the households reported their true needs and obtained water according to their true needs, that would produce the arguably the "best" allocations of the water. Can a mechanism induce the participants to report their preferences truthfully? The term of art for such a mechanism is that it is *incentive compatible*, meaning that the incentives of the individuals are compatible with social goals.

When a mechanism is incentive compatible, in addition to benefits in the quality of the resource allocation, it also yields gains in terms of reducing the cognitive load on the parties. Each party need only act according to its true preferences without regard to what others might be doing: it can ignore subtle strategies and others' decisions, producing simpler demands for knowing others' preferences and reasoning about them. Such a cognitive advantage yields benefits also in terms of perceived fairness, as has been studied for public school assignments for children in Boston Pathak [2011]. Often, it can happen that a mechanism is so complex that people who are more knowledgeable or better connected can figure out cleverer strategies than others and thus have an unfair edge over others. Those who suffer from such arrangements may feel that society is unfair to them and may feel coerced into accepting poorer outcomes than they feel they deserve.

5.4.1 Analyzing the Vickrey Auction

Consider the Vickrey auction with respect to incentive compatibility. As before, we begin by assuming that the (buy) bids are autonomously set by the bidders. The interesting result is that despite their autonomy, bidders would feel "compelled" to reveal their private valuations. Specifically, a rationally dominant strategy is for each bidder to set his or her buy bids to equal his or her private valuation for the item under consideration.

The Vickrey auction is defined for a single seller. Consider two buyers A_1 and A_2, respectively, with private valuations v_1 and v_2, bidding b_1 and b_2.

If $b_1 > b_2$, A_1 wins and pays b_2. A_1's utility in that case is $v_1 - b_2$. This amount is positive or negative depending upon whether A_1's valuation is greater or smaller than the bid placed by the other player. If $b_1 < b_2$, A_1 loses the auction: in this case, A_1's utility is zero, assuming no bidding costs.

Now let us consider the utility to bidder A_1 in the case that A_1 wins. This is $(v_1 - b_2)$ and A_1 should make sure that if A_1 wins, the utility is positive. However, $(v_1 - b_2) > 0$ holds if and only if $v_1 > b_2$. That is, if $v_1 > b_2$, A_1 would benefit by maximizing the probability that $b_1 > b_2$, so that A_1 would win. Given $v_1 > b_2$, the probability is maximized when $b_1 \geq v_1$.

In the same vein, when $v_1 < b_2$, A_1 should minimize the probability of winning the auction—because winning would lead to a net loss. A_1 can do so when $b_1 \leq v_1$. However, A_1 does not know what b_2 is going to be and, specifically, whether b_2 is above or below v_1. Thus, A_1 can choose $b_1 = v_1$, which provides the best outcome in both of the above cases.

That is, underbidding would merely mean that the bidder is likelier to lose, but would pay the same price if it does win. And, overbidding merely mean the bidder is likelier to win but would pay the same price if it does win. Thus, there is no possible gain from overbidding or underbidding—only an increased risk of taking a loss or missing an opportunity to make a gain.

In other words, setting the bid equal to valuation is the best strategy—that is, it dominates every other strategy for determining a bid magnitude knowing the private valuation. Even if the bidder were to know about the other bidders' valuations and bidding strategies, it could do no better than bidding its true valuation.

The general principle demonstrated by the Vickrey auction is that the payoff to a player (given its private valuation) depends only upon the bids of others, not upon its own bid. It is this reliance upon others' actions that grants a certain level of "objectivity" to a method and thereby leads to the possibility of incentive compatibility. Assuming that truthfulness is acceptable at all, then it would be the best policy.

5.4.2 Dominant Strategies

An autonomous party can choose what *strategy* to adopt. A strategy in this setting is a way in which a decision maker arrives at a decision. Specifically, a decision pertains to how much the decision maker should bid for an item given his or her private valuation for that item.

In the setting of auctions, the payoff depends upon how the auction turns out. For those parties who do not trade, we assume the payoff is zero, assuming here that there is no cost merely to participating in the auction, although in practice there exist auctions where there is an entry fee. For the seller, the payoff is the price the seller receives minus the seller's private valuation. For the buyer, the payoff is the buyer's private valuation minus the price the buyer pays.

The payoff that one party receives depends upon both his or her strategy and the strategies of the parties in the environment. Thus, the choice of a strategy can depend upon the concurrent choices of strategies by other entities.

Given all possible strategies that a party may adopt, a *dominant strategy* is one of those that has would produce no worse an outcome for the party regardless of whatever strategies are adopted by other parties. A dominant strategy in general may not exist but when it does the party could simply adopt any of its dominant strategies and not do any worse than otherwise.

Let us consider the first-price auction where buyers bid to purchase an item. The payoff of the winner from such an auction depends upon the winner's private valuation minus the price the winner pays (which is what the winner bid). However, the winner would win even if he or she bid any greater amount and any smaller amount that was at least marginally greater than the second highest bid. To obtain a positive payoff, a buyer should bid a smaller amount than his or her valuation. However, the best amount to bid would depend upon the buyer's estimation of the other bids.

Suppose a buyer with a valuation of \$9 for an item has definite knowledge that the second highest bid would be \$5. The buyer should thus bid \$5.01, thereby assuring a win at the least price to be paid. When buyer has uncertain knowledge, it would choose the price in a more complicated way but nevertheless dependent upon the bids of others.

Now, instead, consider the case of a Vickrey auction. Suppose again that Alice is the winner. Let's suppose her private valuation for the item is \$9. Suppose the highest bid of everyone except Alice is \$7. In this situation, let us consider some alternative bids from Alice.

- Alice bids an amount equal to her private valuation \$9—she pays \$7 and gains a payoff of $2.

- Alice bids $10—she would still net a payoff of $2 (her private valuation minus the price she pays).
- Alice bids $7.01—she would still net a payoff of $2 (her private valuation minus the price she pays).
- Alice bids $6.99—she would still net a payoff of $0 since she would miss out on the transaction.

Given Alice's private valuation of $9 and the highest bid of everyone except Alice being $7, the last of the above alternatives is clearly undesirable for her: she passes an opportunity to make a gain of $2. The other alternatives, however, are equally as good for Alice because each nets her $2. In other words, as long as Alice bids greater than $7, she would do the best she can.

Now consider a situation where Alice's private valuation is $6 whereas the highest of everyone else's bid is $7. Again, in this situation, let us consider some alternative bids from Alice.

- Alice bids $5.99—she would still net a payoff of $0 since she would not miss out on the transaction.
- Alice bids an amount equal to her private valuation $6—she does not win the auction and gains a payoff of zero.
- Alice bids $7.01—she wins the auction and nets a loss of $1 (her private valuation of $6 minus the $7, which is the price she pays).

Given Alice's private valuation of $6 and the highest bid of everyone except Alice being $7, the last of the above alternatives is clearly undesirable for her: she loses a sure $1. The other alternatives, however, are equally as good for Alice because each nets her zero: the best she can hope for. In other words, as long as Alice bids no greater than $6, she would do the best she can.

In other words, when Alice's valuation is greater than everyone else's she gains by bidding more than everyone else. And, when her valuation is lower than the highest of everyone else, she avoids loss by bidding less than the highest of everyone else. That is, Alice's bid should be greater than everyone else's exactly when Alice's private valuation is greater than everyone else's bid. The only way for Alice to ensure this is by bidding an amount equal to her private valuation. That is, under the Vickrey auction, the dominant strategy for a *buyer* is bidding according to its true value.

The deeper motivation is that Alice's payoff is not improved by her strategy because her payoff depends upon her private valuation and *everyone else's* strategy. Alice can only vary the extent to which she participates in the auction by increasing or decreasing her bid relative to her true valuation.

In general, a reliable way to induce incentive compatibility is by making the payoffs to each party independent of that party's revelation. This makes truth telling a dominant strategy for each party.

For the same reason as the above, under first-price auctions, the dominant strategy for a *seller* is to bid its true value. In the same spirit, for single-unit sellers, the M^{th} price auction induces truthfulness. Likewise, for single-unit buyers, the $(M+1)^{st}$ price auction induces truthfulness. The argument we made above for Vickrey auctions fails for multiunit buyers because a buyer may artificially lower some bids to lower the price for other bids.

5.4.3 The Vickrey-Clarke-Groves Mechanism

In general, as we shall see in Section **??**, such mechanisms come with a heavy "cost" in terms of not being able to achieve desirable features. To get warmed up, let's consider a case where the cost in question is purely monetary. This is the notion underlying the Vickrey-Clarke-Groves (VCG) mechanism **VCG** [**VCG**].

The VCG mechanism is based on the idea that you, as a player, can potentially achieve any resource allocation that you want merely by reporting a sufficiently strong need (that is, high valuation) for the resource in question. However, the need you express creates a liability (a "tax") on you that you must discharge. If you are truthful, fine, you pay for what you truly need. If you lie, then you still pay for what you get but if you don't need all of what you get, then you are paying a lot more than you ought to have by remaining truthful. In other words, if you exaggerate your need, you would pay an similarly exaggerated tax to the extent that you would have been better off speaking the truth to begin with.

The VCG mechanism computes the necessary payments from each player based on the *societal cost* of that player's preferences. The intuition behind the societal cost is that it is the loss in social welfare that accrues because of that player's participation in the society. In our water example, if Alice consumes 100 gallons of water a day, that is 100 gallons of water that her neighbors don't get. Similarly, each of her neighbors faces the same situation with respect to everyone else.

5.4.4 Desirable Properties of Markets

Putting the above ideas together, we can identify the following important properties of markets.

- Efficient: it should ensure that the party who values a good the most obtains that good. This is easy enough to ensure by trade. For example, if Alice has a good but values it less than Bob, then the market should support Alice and Bob trading. Efficiency means that a market achieves a states wherein no further gains are possible from trade (players who value goods most get them): i.e., Pareto optimal.

- Truthful: a market encourages participants to bid their true valuations. That is, players would optimize their expected utility by bidding their true valuations. A more general concept here is of *incentive compatibility*, though it is used synonymously with truthfulness in our setting. The idea of incentive compatibility is that the incentives of the individuals are compatible with maximizing social welfare, that is, achieving the most prosocial outcomes. Where resources are allocated through bids the determination of what enhances social welfare most depends upon the true preferences of the participants and thus incentive compatibility reduces to getting the participants to reveal their preferences truthfully.

- Noncoercive: a market preserves the autonomy of any rational individual. That is, no participant is worse off for participating in the market than otherwise. This is synonymous with the market respecting *individually rationality* of the participants.

- Budget balanced: the market neither takes money from the participants nor gives them money. That is, Σ payment = Σ revenue: the seller receives what the buyer pays. In a weaker sense, we can state that there is no subsidy from that market.

Can all of the above properties be satisfied? Here is an interesting argument that that they cannot be satisfied, based on a famous result by Myerson and Satterthwaite. Suppose we have our usual two traders, Alice and Bob, participate in a market. Alice is looking to sell the good that Bob is looking to buy and their valuations come from overlapping distributions. Suppose Alice places a sell bid s and Bob a buy bid of b. The trade can proceed only if Bob bids more than Alice, that is, $b > s$.

If the trade completes, Alice obtains the price p_s and Bob pays the price p_b, these prices to be determined based on the rules of the market. To ensure truthfulness, the market must ensure p_s is independent of what Alice bids and p_b is independent of what Bob bids. That is $p_s = b$ and $p_b = s$. Or, Alice receives b whereas Bob pays s. In other words, this means Alice receives more than Bob pays. Therefore, we have a deficit in the market.

Myerson and Satterthwaite's result means that the market must either subsidize the participants or else we must give up one of the other desirable properties. A subsidy can be practical only if there is some other actor (such as the government) that can sweeten the pot for the participants. In general, one of the other properties would fail. In typical settings, truth telling is the property that we sacrifice.

An alternative design called the *dual-price auction*, due to Preston McAfee [1992], proceeds as follows. Its idea is to determine a price p somewhere in the equilibrium range, that is, between Mth to (M+1)st. Specifically, this price could be the midpoint of the range. The dual-price auction includes all of the possible trades at the dual price but one. Specifically, it omits any trade involving the lowest buyer at or above Mth and the highest seller at or below (M+1)st.

In light of Myerson and Satterthwaite's result, the dual-price auction must violate one of the four properties we saw above. Which one? Clearly, it satisfies budget balance since each trade happens at p: each buyer pays and each seller receives p. It also satisfies truth telling because each trade occurs at a price that is not dependent upon the bids of the trading partners. Each party who trades as a buyer pays a smaller price than his or her bid and each party who trades as a seller receives a larger price than his or her bid. That is, the market is noncoercive: each participant benefits by participating.

However, the dual-price auction loses out on efficiency. By its construction, there is at least one pair of buyer and seller who do not trade but could each gain from the transaction. In essence, we sacrifice these buyers and sellers to produce the information to use for everyone else. However, among all the possible trades that it could discard, the auction discards the least "valuable" trade—least valuable because the gap in the valuations of the omitted pair of buyer and seller is the least of any potentially trading buyer and seller.

5.5 Rationality

Rationality is a basis for understanding interactions among autonomous parties. However, what rationality itself is is far from settled. The classic conception of a rational person is one who is open to reason and not dogmatic: that is, someone who can be persuaded and who can persuade others. In recent times, a more popular conception of a rational person is one who acts in such a manner as to maximize his or her utility given his beliefs at the time of decision. This latter, economic, notion of rationality is the one we discuss here.

To understand economic rationality, let us consider first a space of alternatives or outcomes. Each party a has some ordinal (i.e., sorted) preferences over the alternatives, captured by a binary relation, $\succ_a$. Here, the subscript reflects that it is a whose preferences we are talking about—we omit the subscript when it is clear whom we are talking about. Importantly, the preference ordering $\succ$ is a strict ordering. First, it is asymmetric: If $p \succ q$ then definitely not the reverse, that is, $q \not\succ p$. Second, it is transitive: If $p \succ q$ and $q \succ r$, then $p \succ r$. Of course, $\succ$ is not total, meaning that there are some pairs of alternatives p and q, for which $p \not\succ q$ and $q \not\succ p$. In such a case, we say that the party in question is indifferent between p and q and write it using another binary relation, $\sim$. Specifically, $p \sim q$ means $p \not\succ q$ and $q \not\succ p$.

As in the vernacular, a lottery is a probability distribution over two or more outcomes or alternatives. As usual, the probabilities of the outcomes add up to 1.

A long-established methodology for "eliciting" or inferring the preferences of rational parties is to ask them to choose between lotteries. The odds a person is willing to accept for various outcomes or the odds at which the payoff from the outcomes balances provides information about the valuations a person attaches to various outcomes.

A lottery summarizes the probabilities of occurrence of potential outcomes along with their valuations. For example, if Alice and Bob toss a fair coin to decide who will pay whom a dollar, we can write this as a lottery from Alice's perspective: [0.5: $1; 0.5: -$1]. With 50% chance she gains a dollar; with a 50% chance she loses a dollar. Likewise, if Charlie buys a $10 dollar raffle ticket where he has a one in thousand chance of winning a car he values at $25,000, we can express this raffle as a lottery for him with a 99.9% chance of losing $10 and a 0.1% chance of gaining $24,990. In formal terms, we can write this lottery as [0.001: $24,990; 0.999:-$10].

We define the expected utility or value of a lottery as a scalar (i.e., in monetary terms) as the sum of the payoffs of the different outcomes weighted by probability. For example, if the utility of a car is $25,000, the (expected) value of [0.001: car-$10; 0.999:-$10] equals (0.001*$24,990) + (0.999*(-$10)), which reduces to $15.

We can use lotteries such as the above to discern someone's preferences. Suppose we are trying to figure out how much Alice values a card being raffled. If Alice accepts a lottery (that is, enters into a raffle) for the car, then knowing the probability of her winning and the price she is paying, we can figure out the break-even point for her valuation of the car. For example, consider the lottery [0.01: car-\$10; 0.99:-\$10]. Supposing the value of a car as a sure thing to be c, we can determine the expected value of this lottery as $0.01(c-10)+0.99(-10) = 0.01c-10$. This value is positive only if $c \geq \$1{,}000$. Likewise, if Alice accepts the lottery [0.001: car-\$10; 0.999: -\$10], we can infer that Alice's valuation for the car is at least as a high as \$10,000.

5.5.1 Lotteries and Preferences

The so-called *neoclassical* economic theory makes some assumptions that help establish relationship between lotteries and expected utility.

The notation $\succ$ expresses a player's *preference* relation. That is, $A \succ B$ means that our player prefers outcome A to outcome B. The idea of preferences is that they are qualitative and conceptually unique to each party. It is common to assume that the preference relation is *irreflexive* (if Alice prefers A to B, then she doesn't prefer B to A) and *transitive* (if Alice prefers A to B and B to C, then she prefers A to C). It is customary to define an *indifference* relation, notated as $\sim$, which captures the pairs of outcomes that the player is indifferent between. In general, from a mathematical standpoint, given any two outcomes, a player may neither prefer one to the other nor be indifferent between them.

Another assumption is that a player's preferences are complete, meaning that for any two outcomes, the player prefers one to the other or is indifferent between them. That is, Alice is indifferent between A and B provided she nonstrictly prefers each outcome to the other. In this case, indifference is an equivalence relation. We would ordinarily expect indifference to be *reflexive* (that is, $A \sim A$) and *symmetric* (that is, $A \sim B$ implies $B \sim A$). This assumption adds in *transitivity* (that is, $A \sim B$ and $B \sim C$ implies $A \sim C$).

Based on the foregoing, we can capture some important requirements on preferences and indifference. First, indifferent outcomes are interchangeable or intersubstitutable. Specifically, $A \sim B$, then $[p : A; (1-p) : C] \sim [p : B; (1-p) : C]$. That is, if Alice is indifferent between A and B, then she is indifferent between any two lotteries that differ only with respect to A and B.

Second, a rational player would always want more of what he or she prefers more at the cost of what he or she prefers less. Specifically, if $A \succ B$ and $p > q$, then $[p : A; (1-p) : B] \succ [q : A; (1-q) : B]$. That is, increasing the probability of A over B makes things better for the player who likes A more than B.

Third, the ordinary rules of probability apply to reduce compound or nest lotteries to simple or flat lotteries. That is, $[p : [q : A; 1-q : B]; 1-p : C] = [pq : A; p-pq :$

$B; 1-p:C]$.

A fifth assumption, called *continuity of preferences*, is a somewhat technical one. The intuition here is that a player should be able to price alternatives in terms of each other. For example, suppose Alice prefers icecream to yogurt and yogurt to cookies. Then, according to this assumption Alice would be indifferent between a sure thing of yogurt and some probability mix of icecream and cookies.

5.5.2 Utility Functions

Under the above assumptions, we can map the preferences of a player to a *utility function*, a numeric measure of the gain the player expects from each outcome. That is, a player's utility function—there's exactly one per player—map each alternative (outcome) to a scalar (real number). We can think of this measure as money, although utility is a broader concept than money.

In notation, we write $U : \{\text{alternatives}\} \mapsto \mathscr{R}$. For players with irreflexive, transitive, complete, continuous preferences, there is a utility function U such that these properties hols:

- $U(A) > U(B)$ implies $A \succ B$
- $U(A) = U(B)$ implies $A \sim B$
- $U([p:A; 1-p:C]) = p \times U(A) + (1-p) \times U(C)$ (weighted sum of utilities).

According to the naïve theory, two lotteries with the same expected payoff would have equal utility to a rational party. But, of course, we know that is not so. People are often faced with options with unpredictable outcomes in their business and personal lives and often cope with such unpredictability through additional mechanisms. For example, insurance is a means by which people limit how negative a a negative outcome can be. But assuming the insurance provider is rational and thus profit seeking, insurance inevitably has a negative expected value for the purchaser: on average you would get less than you put in or else the insurance provider would not exist. Yet, buying insurance is far from irrational. The missing element is *risk*: it is rational to pay a premium for reducing the risk one faces.

The flip side of insurance is buying a lottery ticket wherein the gambler would on average lose money. Lotteries, especially state-run lotteries, are sometimes described as a tax on the stupid. Here too, we must not be too hasty to judge. A lottery ticket gambler may figure that the likely loss of the price of a ticket is of far smaller consequence to his economic well being than the unlikely gain from a win.

The economic theory of rationality accommodates risk by modeling it as unpredictability. Let's compare two lotteries in terms of the payoffs of their outcomes (not the outcomes directly): *(a)* $L_1 = [1 : x]$; and *(b)* $L_2 = [p : y; 1 - p : z]$ Suppose $x = py + (1 - p)z$. That is, L_1 and L_2 have the same expected payoff, where L_1 is the sure thing and L_2 is the variable case. Given this, a player's preferences reflect its attitude to risk. A player is respectively risk averse, risk neutral, or risk seeking depending on whether he or she assigns a greater, equal, or smaller utility to the sure thing compared the the variable lottery.

It is customary to make some significant simplifying assumptions. *(a)* The participants are risk neutral.; *(b)* The participants are willing to trade money for any of their resources at a price independent of how much money they already have; and *(c)* The participants know their valuations, which are independent and private

5.5.3 Beyond Simple Utility

Expected utility theory is not predictive of human decision making. In particular, factors besides expected payoff and risk are relevant in real life. For example, how people value a sum of money depends upon factors such as their current wealth: a dollar is worth more to a pauper than to a billionaire. Likewise, \$10 discount for a \$20 toy should be worth the same as a \$10 discount for a \$20,000 car. However, nearly everyone would consider a 50% discount as much more significant than a 0.05% discount. In traditional theory, this is treated as the observation that utility increases with increasing sums of money (as one would expect) but increases at an ever slower rate. That is, a dollar is worth more to someone who has only a hundred dollars than to someone who has a million. Often, utility is treated as the logarithm of the amount one has, the logarithm being a slowly increasing function. For example, $\log(101) - \log(100)$ is a lot larger than $\log(1{,}000{,}001) - \log(1{,}000{,}000)$. In case you are wondering, taking logs to base 10, the former is about four thousandths (0.004,321,373,782,643) whereas as the latter is four millionths (0.000,000,434,294,265).

5.6 Prediction Markets

A prediction market resembles a typical capital (stock) markets in that it is based on what is termed a *continuous double auction*. The motivation both in prediction markets and the capital markers is to accomplish dynamic pricing.

A buy bid is an upper bound: in essence, a buyer declares that he or she is willing to pay up to the bid amount to buy the security. A sell bid is a lower bound: in essence, a seller declares that he or she is willing to accept no less than the bid amount to sell the security.

A double auction is so called because it involves multiple sellers and buyers, potentially with multiple sell and buy bids each. Further, it is called "continuous" or continual because it clears continually. The moment a buyer and seller agree on a price, the deal is done. The matching bids are taken out of the market. Notice that it is quite possible that a moment later a better price may come along, but it will too late then.

In the stock market jargon, *(a)* a *bid* quote is what a seller needs to offer to form a match; *(b)* an *ask* quote what a buyer needs to offer to form a match; and *(c)* the *bid-ask spread* represents the difference in the best (lowest) price a seller would accept and the best (highest) price a buyer would offer.

5.6.1 Potential Pitfalls of Markets

Markets and economic methods in general work on the assumption that the parties are enticed by the positive incentives and discouraged by the disincentives the economic methods offer to the participants. However, this assumption can easily fail.

In commonsense settings, we know about side bets. A gambler may place a side bet in addition to a main bet. A side bet creates a different (additional) payoff for one of the outcomes of the main bet, thereby altering the overall payoff for the gambler. Thus a gambler can appear to be making an irrational choice with respect to the main bet while truly making a rational choice overall. When we don't know about such side bets, therefore, we cannot predict how a rational gambler would behave.

Something similar can happen in prediction markets when one or more of the participants are motivated not by the prices in the market but by other outcomes. That is, unbeknownst to the other players, such players are in essence playing a different game from what everyone else imagines is going on.

A specific example of this arose in one of the prediction markets run by a company called Intrade. This market concerned the 2012 Presidential election in the US wherein the security corresponded to a future event ("Obama wins" or "Romney wins"). Ideally, the price of the security would capture the collective estimate of the probability that the event will occur. It turns out, as Rothschild and Sethi [2015] report, that the Romney wins security was persistently over-valued, suggesting that someone was putting money into that security against their myopic economic self interest as a participant in the prediction market.

A natural explanation is that if a participant would like to bias the opinion of the populace or garner free publicity for his or her favorite candidate (or, more generally, a desired prediction), the participant could purposefully lose money in the market. To someone with a large enough endowment, the loss of funds in the market may be more than offset by the gains in the external world.

In effect, such a player has placed side bets that have changed their overall incentives enough that taking apparently irrational risks in one setting is in fact rational overall.

5.7 Directions

Even in the early days of modern economic theory, researchers recognized the limitations of utility theory as a predictive theory of human behavior. However, these limitations were considered minor or irrelevant in that the economists presented their models as ideals with the backdrop sentiment that people who do not act according to economic theory are irrational and should try to do better.

There are limits of both rationality and our understanding of it. Chapters 9 addresses how the traditional assumptions of economics have played out in light of psychological findings and discussions among others, Tversky and Kahneman's Prospect Theory.

In the 1970s, psychologists began to revisit the assumptions in a systematic way. This has led to a body of work, whose most prominent exponents are Tversky and Kahneman [1981], that largely overturns the assumptions of traditional economics. We return to these new theories in Chapters 9.

Chapter 6 introduces some elements of *social choice theory*.

References

Adam Davidson. 2013. *How Much is Michael Bolton Worth to You?* `http://www.nytimes.com/2013/06/09/magazine/the-secret-science-of-scalping-tickets.html`, June 9.

Iron Maiden Take up Fight against Online Ticket Scalpers. 2015. `http://themusic.com.au/news/all/2015/02/05/iron-maiden-take-up-fight-against-online-ticket-scalpers`, February 5.

R. Preston McAfee. 1992. "A Dominant Strategy Double Auction." *Journal of Economic Theory* 56, no. 2 (April): 434–450.

Paul Milgrom. 1989. “Auctions and Bidding: A Primer.” *Journal of Economic Perspectives* 3, no. 3 (Summer): 3–22.

Parag A. Pathak. 2011. “The Mechanism Design Approach to Student Assignment.” *Annual Review of Economics* 3, no. 1 (March): 513–536.

William Poundstone. 1992. *Prisoner’s Dilemma.* New York: Doubleday.

David M. Rothschild and Rajiv Sethi. 2015. *Wishful Thinking, Manipulation, and the Wisdom of Crowds: Evidence from a Political Betting Market.* TR 2322420. `http://ssrn.com/abstract=2322420`. Rochester, New York: Social Science Research Network, June.

Martin Shubik. 1971. “The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation.” *Journal of Conflict Resolution* 15, no. 1 (March): 109–111.

Amos Tversky and Daniel Kahneman. 1981. “The Framing of Decisions and the Psychology of Choice.” *Science* 211, no. 4481 (January): 453–458.

