

DYNAMIC AUCTIONS

DIRK BERGEMANN
Department of Economics, Yale
University, New Haven,
Connecticut

MAHER SAID
Microsoft Research New
England, Cambridge,
Massachusetts

Olin Business School,
Washington University in
St. Louis, St. Louis, Missouri

Many important economic problems have been studied with the tools of auction theory and mechanism design more generally. Much of the literature, however, studies a static, one-time decision. In many problems of interest, more than a single decision needs to be made; rather, a sequence of decisions needs to be made. These decisions often depend crucially on the dynamic aspects of the environment. For example, in the classical airline revenue management problem, an airline must decide how to price seats on a flight in response to its changing inventory as well as to the evolution of the customer base. A search engine must choose how to allocate its advertising inventory in response to changing search queries and advertiser budgets. Network capacity (either bandwidth or computational resources) needs to be dynamically reallocated in response to the arrival of new computational tasks of varying priority. Private and public procurement agencies must decide how to assign new contracts in response to the changing experiences, competencies, and availabilities of contractors.

The theory of auctions and mechanism design has been extremely successful in tackling a wide variety of problems. Unfortunately, solutions to static problems often do not translate directly to dynamic settings. Consider, for instance, the standard second-price auction for a single object with private values. Buyers are asked to submit bids for

the good, which is then allocated to the highest bidder. This winning bidder then pays the second-highest of the submitted bids. As first demonstrated by Vickrey [1], it is a (weakly) dominant strategy for buyers to bid their true values. Thus, the object is allocated efficiently: the buyer who values the object the most receives it.

In a dynamic setting, however, the second-price auction need not retain its desirable properties. The following example is illustrative. Suppose that there are two objects to be allocated: use of a production facility today and use of that same facility tomorrow. Today, two perfectly patient buyers wish to make use of the facility, the first valuing that use at \$100 and the second valuing it at \$75. A third buyer who will arrive tomorrow values use of the facility at \$50. If a second-price auction is used in each period and the buyers bid truthfully, then the first buyer will win in the first period at a price of \$75, and the second buyer will win in the second period at a price of \$50. On the other hand, suppose that the first buyer instead bids \$60 in the first period and truthfully in the second. She will then lose in the first period, but win in the second and pay a price of \$50, below what she otherwise would have paid. Thus, bidding one's value need not be an optimal strategy in a sequence of second-price auctions.

Notice that in the above example, despite the lack of truthful behavior by the bidders, the outcome is still efficient. As we will discuss later, this need not always be the case. In a more general setting when buyers are uncertain about the arrival time and willingness-to-pay of each of their competitors, Said [2] demonstrates that the second-price sealed-bid auction is often incapable of yielding an efficient outcome, even when bidders are forward-looking and fully rational. Thus, in order to achieve desirable outcomes in dynamic settings, we must move beyond tools that are best suited to static environments.

This is precisely the aim of the recent (and growing) literature on dynamic mechanism

design. We survey this literature, categorizing it by the nature of the dynamics involved. In particular, we examine two strands of literature: one in which the population of agents changes over time, but their private information is fixed; and the other in which the population of agents is fixed, but their private information changes over time. We will further divide these two categories into those papers whose aim is achieving efficient outcomes and those whose aim is achieving revenue-maximizing outcomes.

A DYNAMIC POPULATION WITH FIXED INFORMATION

We begin by exploring the possibility of implementing desirable outcomes when faced with a dynamic population whose private information is fixed. In the course of our discussion, we will examine some of the problems that arise in such settings and how they differ from those found in static, one-shot models. Note that all the models and papers we discuss assume private values.

Maximizing Social Welfare

Efficiency in the Face of Arrival and Departure Dynamics. Parkes and Singh [3] develop a framework for studying sequential allocation problems in a dynamic population setting. In their world, various subsets of agents are present at various times. In particular, each agent has an arrival time and a departure time, and has a utility function for decisions made while she is present. What makes the problem especially difficult is that the decision maker (and mechanism designer) does not know any of this information—the designer must incentivize the agents to reveal that information.

Despite this difficulty, Parkes and Singh are able to construct a mechanism (the “online VCG” mechanism) which generalizes the Vickrey-Clarke-Groves mechanism from static environments and implements the efficient outcome. Their mechanism is a direct revelation mechanism in which buyers report their private information upon arrival to the mechanism. Given these reports, the mechanism carries out an efficient policy.

In order to enable the truthful revelation of private information, the mechanism chooses transfers that align agents’ incentives with those of a benevolent social planner. Similar to how the static VCG mechanism sets transfers equal to the externality that an agent imposes on other agents, the online VCG mechanism sets transfers equal to the *expected* externality that an agent imposes. This expected externality takes into account each agent’s impact on the other agents currently present *as well as* on agents that may arrive in the future.

We should note, however, that *unlike* the static VCG mechanism, this truthful reporting is *not* a dominant strategy, as the calculation of expected externalities depends upon the truthful reporting of all other agents. If, instead, an agent believes that agents in the future are not reporting truthfully, it may be in her best interest to misreport her true type.

Efficient Sequential Assignment with Impatient Buyers. In related work, Gershkov and Moldovanu [4] examine the allocation of a finite set of heterogeneous durable goods to a dynamic population of randomly arriving buyers. Crucially, the buyers in this setting are impatient—they wish to purchase an object immediately upon their arrival on the market. Moreover, these buyers share a common ranking of the various objects. In this setting, objects are durable while buyers are impatient. This implies that the relevant trade-off is between allocating an object to a buyer at the current time and the option of assigning it to a buyer in the future who might value it more. In a complete information setting, the efficient dynamic allocation policy was first characterized by Albright [5]. This policy prescribes an ordered partition of the set of possible valuations at each possible time. Agents with values in the highest element of the partition receive the best available object; agents in the second-highest element receive the second-best remaining object, and so on. These intervals may depend upon the time of an agent’s arrival or on the set of remaining objects.

Gershkov and Moldovanu [4] show that this efficient policy is, in fact, implementable in the presence of incomplete information

about the valuations of arriving buyers. Since buyers are impatient, there is no loss of generality in considering direct revelation mechanisms where each buyer reports her type to the mechanism designer upon her arrival to the market. Each buyer is charged the *expected* externality that he/she imposes on future agents, where expectations are taken with respect to the arrival process and the valuations of future agents. This implies that each agent's net utility is exactly equal to his/her expected marginal contribution to the social welfare, conditional on the information available at the time of her arrival—essentially, the efficient mechanism is a dynamic Vickrey-Clarke-Groves mechanism like the Parkes and Singh [3] online VCG mechanism discussed above.

Efficient Auctions with Dynamic Populations. In a closely related recent work, Said [2,6] examines the allocation of a sequence of indivisible goods to a dynamic population of buyers. In this setting, patient buyers demand a single unit of perishable homogeneous goods (such as computational time on a central computing facility, or capacity in a production line). These buyers arrive at random times to the market, and have private valuations, which are drawn independently and identically from a given distribution.

While a variant of the online VCG mechanism of Parkes and Singh [3] may be used in order to achieve an efficient allocation of objects to buyers, these mechanisms are direct revelation mechanisms, requiring buyers to report their values to the mechanism upon their arrival to the market. In practice, however, direct revelation mechanisms may be difficult to implement or undesirable due to privacy or complexity concerns.

Therefore, *indirect* auction mechanisms are often useful. In standard static settings, the canonical auction for efficiently allocating goods to buyers is the sealed-bid second-price auction. This mechanism is the auction analogue of the Vickrey-Clarke-Groves mechanism, allocating the object to the buyer who submits the highest bid and charging her a price equal to the second-highest bid. Said [2] shows that, when selling a sequence of objects, one in each period, to a stream

of buyers who arrive at random times, the sealed-bid second-price auction is no longer efficient.

In a sequential auction, buyers have an “option value” associated with losing in a particular auction, as losing bidders have the possibility of winning in a future auction. The value of this option to a particular buyer depends on his/her expectations of the prices he/she may have to pay in the future. This is, of course, dependent upon the private information and values of other competitors. Thus, despite the fact that buyers' values for an object do not depend upon the information of their competitors, the dynamics of a sequential auction market *generate interdependence* in (option) values.

This interdependence implies that a standard second-price sealed-bid auction does not reveal sufficient information for the determination of buyers' option values. In contrast, the ascending (or English) auction is a simple *open* auction format that does allow for the gradual revelation of buyers' private information. As buyers drop out of each auction, they indirectly reveal their private information to their competitors, who are then able to condition their current-period bids on this information. Moreover, this process is repeated in every auction in every period. This allows newly arrived buyers to learn about their competitors without being privy to the detailed events of the previous periods, thereby allowing for an efficient outcome.

Maximizing Revenue

Before discussing revenue-maximizing dynamic mechanisms and auctions, it will be helpful to briefly review the (standard) results for static environments. In a static setting, in which all buyers are present simultaneously, Myerson [7] solves the optimal mechanism design problem in the following manner:

1. Incentive compatibility requires that the allocation rule be monotone; that is, buyers who report a higher type must have a greater probability of receiving an object.
2. Given a monotone allocation rule, the corresponding payment rule is pinned

down (up to a constant) via the revenue-equivalence principle. When the goal is maximizing revenues, the constant can be determined via individual rationality constraints.

3. Combining these two observations, the revenue-maximization problem can be transformed to the problem of maximizing “virtual” surplus—the surplus from allocating objects less the information rents that must be paid in order to induce truthful revelation of types.

Thus, instead of maximizing the sum of agents’ values as in the case of efficient mechanism design, Myerson’s insight is that revenue-maximization is equivalent to maximizing the sum of agents’ *virtual* values, where the virtual value of a buyer is given by the difference between his/her true value and the information rents necessary for inducing truth-telling. An optimal allocation rule, therefore, allocates objects to the buyers with the highest nonnegative virtual values—precisely those with the most favorable trade-off between surplus and information rents.

When buyers are *ex ante* symmetric and the distribution of values is such that virtual values are nondecreasing then this optimal allocation rule is monotone in values. Thus, incentive compatibility is achieved “for free” and the optimal auction corresponds to a second-price auction with an appropriately chosen reserve price. If, on the other hand, virtual values are not increasing, then the monotonicity constraint may be satisfied via a procedure termed “ironing” which pools some types of agents and randomizes among them. Thus, in the static setting, increasing virtual values is sufficient for revenue-maximization to correspond to maximizing virtual surplus.

Revenue-Maximization in the Face of Arrival and Departure Dynamics. Returning to the setting of arrival and departure dynamics, Pai and Vohra [8] present what may be thought of as the revenue-maximizing counterpart to the efficient online VCG mechanism of Parkes and Singh [3]. Here, a seller with a fixed, finite supply of a

homogenous good faces a population of potential buyers with unit demand who arrive and depart over the course of a finite time horizon. As in Parkes and Singh, the times at which each agent arrives and departs from the market are her private information, as is her valuation for an object. Agents’ types are assumed to be independent draws from a common distribution. The seller wishes to design an incentive-compatible direct mechanism in order to maximize his revenues.

The challenge then is to extend Myerson’s notion of a virtual value in order to account for the additional privately known arrival and departure times. Moreover, even if such a virtual value is defined, the multidimensional nature of the private information implies that we must also find an appropriate notion of monotonicity in order to guarantee incentive compatibility.

Some care must be used when considering direct revelation mechanisms in this setting with random arrival and departure times. In particular, note that buyers cannot make reports before their arrival or claim to depart after their true departure time (see Green and Laffont [9] for more on restricted mechanisms and the revelation principle). Taking this into account, Pai and Vohra use an appropriate formulation of the revelation principle of Myerson [10] to restrict attention to mechanisms that allocate to buyers, if at all, only in the period of their departure. They use the incentive compatibility constraints to show that the optimal allocation rule must be monotone in valuations, holding entry and exit times fixed. Moreover, incentive compatibility also requires that the allocation rule be monotone in entry and exit times. In particular, a buyer who arrives earlier or departs later should have a greater probability of receiving an object. (Thus, buyers are incentivized to report the greatest possible window of opportunity.)

Making use of the conditions imposed by incentive compatibility, Pai and Vohra [8] are able to arrive at a notion of virtual valuation similar to that of Myerson [7], except that information rents vary across reported arrival and departure times. Thus, in each period, the seller compares the

virtual surplus from allocating an object to a departing agent to the expected virtual surplus of waiting an additional period. A departing agent will receive an object only if the former is greater than the latter. However, unlike the setting of Myerson, increasing virtual values are no longer sufficient for such an allocation rule to be optimal, but instead additional “ironing” may be necessary: the optimal mechanism may need to occasionally withhold an object despite a lower future expected virtual surplus.

Optimal Sequential Assignment with Impatient Buyers. Recall the relationship between Parkes and Singh [3] and Gershkov and Moldovanu [4] discussed earlier. An analogous relationship exists between the work of Pai and Vohra [8] and Gershkov and Moldovanu [11]. This latter work considers the problem of revenue-maximization in a setting with a dynamic population of impatient buyers. As in the paper discussed earlier, a seller has a finite set of heterogeneous durable goods that he wishes to sell to randomly arriving buyers who wish to purchase an object immediately upon their arrival to the market. (In a setting with homogeneous durable goods and patient buyers arriving stochastically, Gallien [12] shows that when the distribution of buyer inter-arrival times has an increasing failure rate, the optimal mechanism allocates goods to buyers only upon their arrival. There, as Board [13] also shows in a related discrete-time setting, the seller may behave as though buyers are impatient.)

As in Pai and Vohra [8], there is a dynamic trade-off that is introduced by the arrival process of buyers. In particular, when considering allocating an object to a particular agent, the seller must consider the opportunity cost of allocating that object to future agents. However, the problem here is (in some ways) simpler as buyers no longer have multidimensional types. Since buyers are impatient, there is no need to consider incentives for revealing arrival or departure times—the only dimension of private information is the private value that buyers have for each object.

Thus, the authors are able to prove that (as in the case of the efficient policy) any incentive compatible policy must take the form of an ordered partition of possible valuations at each point in time. This follows directly from the fact that short-lived agents’ decision problems are static, and so static insights about incentives apply directly. Thus, agents with values in the highest element of the partition receive the best available object; agents in the second-highest element receive the second-best remaining object, and so on.

The next step in the characterization of the revenue-maximizing dynamic mechanism is to determine the revenue generated by any particular incentive compatible policy. The problem faced by the seller upon the arrival of any particular agent can then be translated into a static problem. More specifically, the seller must decide between allocating an object to the newly arrived agent and obtaining the revenue generated by ignoring that agent’s arrival and proceeding with the (given) policy. This is the same as the problem faced by a standard (static) revenue-maximizing monopolist selling a single object, when the cost of that object is given by the salvage value of the inventory (which corresponds to the expected future revenues from the dynamic policy). It is possible, by working backwards from the case of a single-object inventory, to determine the optimal partition of the valuation space for each possible inventory and point in time. The revenue-maximizing prices that correspond to these cutoffs are then derived via revenue-equivalence.

Optimal Auctions with Dynamic Populations. In a complementary setting, Said [2] explores the properties of revenue-maximizing dynamic auctions. Instead of durable goods and impatient buyers, he examines the assignment of a sequence of perishable goods to a population of patient buyers. The buyers arrive to the market at random times and remain until they receive an object. Since their values for the objects are private information, incentives must be provided for information revelation in order to achieve an optimal allocation.

In this setting, the seller and buyers discount the future with a common discount

factor $\delta \in (0, 1)$. This implies that even though objects are assumed to be individual units of a homogeneous good, from the perspective of any individual buyer, they are differentiated products. To make this clear, consider a buyer i with value v_i who is present on the market at some time t . If this buyer receives an object in period t , she receives a payoff of v_i . However, if she anticipates receiving an object in period $t + 1$, then her expected payoff is δv_i . Thus, she does not value the two objects identically.

This alters the appropriate condition for incentive compatibility. Recall that in static settings with single-unit demand, incentive compatibility requires that increasing a buyer's type should increase her probability of receiving an object. In this dynamic setting with single-unit demand, incentive compatibility instead requires that increasing a buyer's type should increase her probability of receiving an object sooner. This characterization of incentive compatibility allows for the derivation of a revenue-equivalence theorem for this setting, which implies that revenue-maximization can be achieved (in the Myersonian tradition) by maximizing the virtual social surplus.

Said [2] shows that applying a variant of the online VCG mechanism discussed earlier to virtual values maximizes revenues. Instead of providing each buyer with a net utility equal to her expected marginal contribution to the social surplus, Said's mechanism leaves each buyer with a net utility equal to her expected marginal contribution to the *virtual* surplus. Moreover, if an indirect mechanism is desired, the earlier discussion on efficient but indirect mechanisms remains relevant—instead of a sequence of second-price sealed-bid auctions with an optimally chosen reserve price, a seller should use a sequence of ascending auctions with that same reserve price.

If objects are durable instead of perishable, the recent work of Board and Skrzypacz [14] provides a characterization of the optimal dynamic auction. As before, the revenue-maximizing direct mechanism applies an efficient mechanism to virtual values. This implies that there is a constant cutoff value below which objects are not

allocated, and the object is allocated to the buyer with the highest valuation exceeding that cutoff. This cutoff is determined by a simple one-period-look-ahead policy which requires the seller to be indifferent between selling to the cutoff type in the current period and waiting an additional period for new buyers to arrive.

With a single durable object to be allocated, these cutoffs are constant in all but the final period, and then drop in that period to the static monopoly price. This implies that buyers either buy immediately upon their arrival or wait for the “fire sale” in the final period. If the seller wishes to use an indirect mechanism in order to implement this optimal policy, Board and Skrzypacz show that a sequence of second-price auctions with deterministically declining reserve prices is optimal. These reserve prices are chosen such that the cutoff type is exactly indifferent between purchasing in the current period and waiting an additional period. The reserve price must be declining so as to compensate this cutoff-type buyer for the losses due to discounting *as well as* for the possibility of new competition arriving in future periods. It follows that the reserve price must fall at a rate faster than the discount rate.

A FIXED POPULATION WITH DYNAMIC INFORMATION

The preceding section examined dynamic auctions and mechanisms for settings in which there is a dynamic population with static, fixed types. A large and important class of problems, however, does not fit within this framework. Consider, for instance, the case of sponsored search, where search engines sell advertising slots alongside “organic” search results. In this setting, advertisers who are interested in a particular keyword may revise their estimates of the value of an advertisement based on their experiences over time. The natural question that arises in such a setting is how to design mechanisms and auctions in order to properly account for the changing private information.

We first consider the case of implementing socially efficient outcomes in environments

with dynamic information. The mechanisms we examine generalize insights from static settings to enable the truthful revelation of changing private information in a dynamic environment. We then turn to the objective of maximizing revenues and discuss the recent work in that area. As in the previous section, we focus on private value environments.

Maximizing Social Welfare

Dynamic Pivot. Bergemann and Välimäki [15] consider the problem of providing incentives for truth-telling and satisfying participation constraints in an efficient mechanism. They consider a general infinite-horizon dynamic model in which participants observe a sequence of private signals over time, and a sequence of decisions must be taken in response to these signals. The distribution of signals may depend on previously observed signals or decisions; however, these signals are independent across agents (conditional on observables).

Bergemann and Välimäki propose the Dynamic Pivot Mechanism. This mechanism is the natural generalization of the static Vickrey-Clarke-Groves mechanism to a dynamic environment. Recall that VCG-like mechanisms incentivize truth-telling by making efficient decisions and choosing transfers such that each agent's payoff is equal to the total social surplus. By varying the transfers to each agent in a way that does not depend on her own report, we obtain an entire set of efficient mechanisms. The "Pivot" mechanism (which is often considered the canonical VCG mechanism) is the member of this class that subtracts from each agent's transfer the surplus that the other agents could have achieved in her absence; that is, it chooses transfers equal to the externalities imposed on the rest of the system so that each agent's net utility is equal to her marginal contribution to the social welfare.

The key insight of Bergemann and Välimäki [15] is that in a setting with dynamic private information, incentives for truth-telling and efficiency may be guaranteed by choosing payments in each period to be equal to the *flow* externalities. More precisely, the externality that an

agent's report in any given period imposes on all other participating agents may be decomposed into a current period effect and a future (expected) effect. By committing to a sequence of transfers equal to the sequence of current-period externalities, the mechanism designer or social planner is able to give each buyer a flow payoff equal to her flow marginal contribution to the social welfare. Therefore, each agent's expected discounted payoff, looking forward from any period, is equal to her *total* marginal contribution to the social welfare (again, starting at that point). This property (which figured prominently in Bergemann and Välimäki [16,17] in the construction of dynamic first-price auctions in a complete information setting) implies that, in each period, all participating agents internalize the impact of their current reports on others, thereby aligning their incentives with those of the efficiency-minded planner.

Cavallo, Parkes, and Singh [18] develop a mechanism similar to the Dynamic Pivot Mechanism in a setting with agents, whose type evolution follows a Markov process. In Cavallo, Parkes, and Singh [19], they extend dynamic VCG to settings in which buyers are "periodically inaccessible" and are unable to make reports (as in the case, for instance, of a network setting where connectivity is sometimes lost). Moreover, Cavallo [20] has shown that the dynamic pivot mechanism is also similar to its static VCG counterpart, in that it achieves greater revenue than any other possible efficient (dynamic) mechanism. Therefore, in settings where the mechanism designer is concerned with both revenue and efficiency, the use of the dynamic pivot mechanism may provide a partial resolution to the conflict among the two.

Efficiency and Budget-Balance. Athey and Segal [21] consider a similar setting to that of Bergemann and Välimäki [15]; however, they are also interested in finding an efficient mechanism that is budget-balanced—that is, a mechanism that requires no external subsidies. In order to do so, Athey and Segal propose the efficient "Team Mechanism." This mechanism induces truthfulness by providing a transfer to each agent, in each

period that is equal to the sum of utilities of all the other agents in that period. Thus, in the same manner as the Dynamic Pivot Mechanism (as well as the static VCG mechanism), such a transfer scheme ensures that each agent is a residual claimant for the total social surplus.

As the Team Mechanism is not budget-balanced, Athey and Segal modify it in a generalization of the classic AGV mechanism of d'Aspremont and Gerard-Varet [22]. The AGV mechanism, in a static environment, permits truth-telling by charging agents transfers equal to the “expected externality,” their reports impose on others. However, the value of these transfers are dependent upon agents’ beliefs, and in a dynamic setting, these beliefs may be evolving over time and may, in fact, be manipulated by other agents. In order to avoid such manipulations, Athey and Segal construct the “Balanced Mechanism,” which takes advantage of the dynamic setting using the *changes* in the expected present value of other agents’ utilities. By doing so, agents continue to internalize the expected externality that they impose upon others. Moreover, these new transfers are not manipulable by other agents, thereby allowing for budget-balance while also implementing efficient outcomes.

One should note that, as with the static AGV mechanism, it may be impossible to both balance the budget and also satisfy participation constraints. Although the Balanced Mechanism satisfies *ex ante* participation constraints (so that all agents wish to participate in the mechanism at the time of contracting), it differs from the Dynamic Pivot Mechanism of Bergemann and Välimäki [15], in that some agents may wish to exit the mechanism after some histories. Athey and Segal [21] show, however, that if the time horizon is infinite and agents are sufficiently patient, then it is possible to satisfy these “periodic” participation constraints.

Maximizing Revenue

Contracting with Dynamic Information. In many environments, buyers learn about their demand over time. In an important contribution, Courty and Li [23] consider optimal

dynamic contracts for such a setting. In their model, buyers have private information about the future distribution of their valuations for the seller’s object. Subsequently, these buyers then privately learn the realization of their value. The monopolist could choose to wait until buyers learn their realized values and then charge the standard monopoly price. Surprisingly, by requiring the buyers to reveal their private information sequentially, the seller is able to extract additional surplus. Instead of “screening” buyers once after all private information is acquired, revenue-maximization requires screening them twice with a set of contingent contracts.

As in a standard static mechanism design problem, revenue-maximization corresponds to allocating to the buyers with the highest virtual values, thereby maximizing virtual surplus. Since buyers’ initial types influence their ultimate value for the good, the concept of a virtual value in this context must take into account this influence. Specifically, the information rents “paid” to buyers in the first period (at the time of contracting) must account for the informativeness of initial types about future types: the more informative the initial type about the future realization, the greater the information rents that must be paid. Since incentive compatibility requires monotonicity in the allocation rule, this implies that distortions away from efficiency are correspondingly increased.

Esó and Szentes [24] consider a closely related setting in which buyers have an initial estimate of their valuation, and the seller is able to control buyers’ acquisition of additional refined information. Specifically, the seller is able to choose whether to allow buyers to receive additional signals that are correlated with their private information. They show that in the optimal revenue-maximizing mechanism, the seller releases all the additional private information that he is able to. Moreover, his revenue in this case is as large as the revenue he would be able to achieve if he could observe this information—quite surprisingly, buyers gain no advantage from this additional information and receive no additional information rents when the seller is able to control

the flow of information. Roughly speaking, by using a more sophisticated dynamic mechanism, the seller pays information rents to buyers in the initial period, in order to incentivize the revelation of their initial private information, but subsequently extracts all remaining potential surplus.

In a particular application, Esó and Szentes [24] show how such a mechanism would operate. Suppose buyers' initial private information is a noisy signal of their value, and that they can then receive additional information that informs them of their actual value. The optimal mechanism in this case is a "handicap auction" in which buyers can purchase an advantage in a second-stage efficient auction. More concretely, each buyer purchases a price premium from the seller (where a smaller premium costs more), and then in the second period (after the revelation of additional information) the buyers engage in a modified second-price sealed-bid auction in which the winner is required to pay her purchased premium in addition to the second-highest bid. The seller is able to extract additional revenue by inducing buyers with higher expected values to purchase smaller handicaps for the second-stage auction.

Dynamic Incentives and Revenue Equivalence. The recent contribution of Pavan, Segal, and Toikka [25] generalizes the results of Myerson [7] and Baron and Besanko [26] to a general dynamic setting. In particular, they characterize incentive compatibility and revenue in multiperiod settings with dynamic private information. The primary challenge in such settings is an element of multidimensionality. Instead of the single dimension of potential misreports possible in a traditional static mechanism design problem, agents may be able to misrepresent multiple "pieces" of information, possibly choosing misreports contingent on her previous private information.

The key step in their analysis is the development of a "dynamic payoff formula" which expresses the derivative of an agent's expected payoff (and hence the seller's revenue) in an incentive compatible mechanism with respect to her private information.

This formula summarizes the local incentive compatibility constraints that a mechanism designer must respect. Pavan, Segal, and Toikka demonstrate the manner in which this dynamic payoff formula relies on incentive compatibility not only in a given period, but also in all future periods; therefore, they are able to capture not just the impact of a change in an agent's private information on her current-period payoffs, but also the influence on future private information.

This machinery yields a dynamic revenue-equivalence theorem: the expected revenue in any dynamic mechanism is determined entirely (up to a constant) by the allocation rule. Moreover, the expected revenue in an incentive compatible mechanism is given exactly by the expected virtual surplus generated by that mechanism, where agents' virtual values (as in Baron and Besanko [26]; Country and Li [23]; and Esó and Szentes [24]) depend upon the informativeness of current-period private information on future private information. Thus, Pavan, Segal, and Toikka [25] provide necessary conditions (and some sufficient conditions) for incentive compatibility and a general methodology for revenue-maximization in dynamic settings.

One limitation of the generality of that work, however, is the difficulty in establishing general sufficient conditions for the incentive compatibility of a mechanism that are easily verified. In a complementary paper, Pavan, Segal, and Toikka [27] use an approach pioneered by Esó and Szentes [24] to represent an agent's future-period types as a function of their first-period type and a random "shock" that is independent of the first-period type. This approach allows the identification (under slightly different assumptions) of conditions for incentive compatibility in infinite-horizon settings.

Despite this, however, sufficient conditions for incentive compatibility often need to be determined on a case-by-case basis for various models. One such determination may be found in Kakade, Lobel, and Nazerzadeh [28], which sets out a model of sponsored search auctions in which advertisers learn about the efficacy of their advertisements over time. Their model combines elements of both private information and statistical

learning. In particular, private information evolves according to a multiarmed bandit process (Deb [29] also considers a similar setting). The authors also impose “separability” conditions on the stochastic process governing the evolution of private information (specifically, that it is a Markov process) along with additive or multiplicative separability of buyers’ utility functions in their initial private information and future private information. These conditions are, in fact, sufficient conditions under which the maximization of virtual surplus by using an efficient algorithm (the Gittins index policy) applied to virtual values maximizes revenue.

In a related setting, Battaglini [30] examines an infinite-horizon model and characterizes the revenue-maximizing long-term contract of a monopolist facing a buyer whose value follows a two-type Markov process. This optimal contract is nonstationary; however, this contract (eventually) converges over time to the efficient supply. A crucial insight (that follows from the work of Pavan, Segal, and Toikka) is that, in an irreducible two-type Markov process, the impact of changing the current-period type vanishes in the long run—the current-period type is almost completely uninformative about types in the distant future. Therefore, distortions away from the efficient contract diminish in the long run, regardless of the degree to which types are persistent.

FURTHER READING

A large body of work has developed in this area. For further reading and examples of dynamic auctions and mechanisms for settings with changing populations and private information, the curious reader may wish to consult, among others, Akan, Ata, and Dana [31]; Aviv and Pazgal [32]; Cav-allo [33]; Doepke and Townsend [34]; Gershkov and Moldovanu [35]; Kittsteiner and Moldovanu [36]; Parkes [37]; Satterthwaite and Shneyerov [38]; Shen and Su [39]; Talluri and van Ryzin [40] [Chapter 6]; and Vulcano, van Ryzin, and Maglaras [41].

REFERENCES

1. Vickrey W. Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 1961;16(1):8–37.
2. Said M. Auctions with dynamic populations: efficiency and revenue maximization. Proceedings of the 1st Conference on Auctions, Market Mechanisms and Their Applications (AMMA 2009); 2009; Boston (MA).
3. Parkes DC, Singh S. An MDP-based approach to online mechanism design. Proceedings of the 17th Annual Conference on Neural Information Processing Systems (NIPS’03). Vancouver, Canada: 2003.
4. Gershkov A, Moldovanu B. Efficient sequential assignment with incomplete information. *Games Econ Behav* 2010;68(1):144–154.
5. Albright SC. Optimal sequential assignments with random arrival times. *Manage Sci* 1974;21(1):60–67.
6. Said M. Information revelation and random entry in sequential ascending auctions. Proceedings of the 9th ACM Conference on Electronic Commerce (EC’08); 2008; Chicago (IL).
7. Myerson RB. Optimal auction design. *Math Oper Res* 1981;6(1):58–73.
8. Pai M, Vohra R. Optimal Dynamic Auctions. Unpublished manuscript, Northwestern University; 2009.
9. Green JR, Laffont J-J. Partially verifiable information and mechanism design. *Rev Econ Stud* 1986;53(3):447–456.
10. Myerson RB. Multistage games with communication. *Econometrica* 1986;54(2):323–358.
11. Gershkov A, Moldovanu B. Dynamic revenue maximization with heterogeneous objects: a mechanism design approach. *Am Econ J Microecon* 2009;1(2):168–198.
12. Gallien J. Dynamic mechanism design for online commerce. *Oper Res* 2006;54(2):291–310.
13. Board S. Durable-goods monopoly with varying demand. *Rev Econ Stud* 2008;75(2):391–413.
14. Board S, Skrzypacz A. Optimal dynamic auctions for durable goods: posted prices and fire-sales. Unpublished manuscript, Stanford University; 2010.
15. Bergemann D, Välimäki J. The dynamic pivot mechanism. *Econometrica* 2010;78(2):771–789.
16. Bergemann D, Välimäki J. Dynamic common agency. *J Econ Theory* 2003;111(1):23–48.

17. Bergemann D, Välimäki J. Dynamic price competition. *J Econ Theory* 2006;127(1):232–263.
18. Cavallo R, Parkes DC, Singh S. Optimal coordinated planning among self-interested agents with private state. Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI'06); 2006; Cambridge (MA). pp. 55–62.
19. Cavallo R, Parkes DC, Singh S. Efficient mechanisms with dynamic populations and dynamic types. Unpublished manuscript, Harvard University; 2009.
20. Cavallo R. Efficiency and redistribution in dynamic mechanism design. Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08); 2008; Chicago (IL).
21. Athey S, Segal I. An efficient dynamic mechanism. Unpublished manuscript, Harvard University; 2007.
22. d'Aspremont C, Gérard-Varet L-A. Incentives and incomplete information. *J Public Econ* 1979;11(1):25–45.
23. Courty P, Li H. Sequential screening. *Rev Econ Stud* 2000;67(4):697–717.
24. Esó P, Szentes B. Optimal information disclosure in auctions and the handicap auction. *Rev Econ Stud* 2007;74(3):705–731.
25. Pavan A, Segal I, Toikka J. Dynamic mechanism design: revenue equivalence, profit maximization and information disclosure. Unpublished manuscript; Northwestern University; 2009.
26. Baron DP, Besanko D. Regulation and information in a continuing relationship. *Inf Econ Policy* 1984;1(3):267–302.
27. Pavan A, Segal I, Toikka J. Infinite-horizon mechanism design: the independent-shock approach. Unpublished manuscript, Northwestern University; 2010.
28. Kakade SM, Lobel I, Nazerzadeh H. An optimal dynamic mechanism for multi-armed bandit processes. Unpublished manuscript, Microsoft Research New England; 2010.
29. Deb R. Optimal contracting of new experience goods. Unpublished manuscript, Yale University; 2008.
30. Battaglini M. Long-term contracting with Markovian consumers. *Am Econ Rev* 2005;95(3):637–658.
31. Akan M, Ata B, Dana J. Revenue management by sequential screening. Unpublished manuscript, Carnegie Mellon University; 2009.
32. Aviv Y, Pazgal A. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manuf Serv Oper Manage* 2008;10(3):339–359.
33. Cavallo R. Mechanism design for dynamic settings. *ACM SIGecom Exchanges* 2009; 8, (2).
34. Doepke M, Townsend RM. Dynamic mechanism design with hidden income and hidden actions. *J Econ Theory* 2006;126(1):235–285.
35. Gershkov A, Moldovanu B. Learning about the future and dynamic efficiency. *Am Econ Rev* 2009;99(4):1576–1587.
36. Kittsteiner T, Moldovanu B. Priority auctions and queue disciplines that depend on processing time. *Manage Sci* 2005;51(2):236–248.
37. Parkes DC. Online mechanisms. In: Nisan Noam, Roughgarden Tim, Tardos Eva, *et al.*, editors. *Algorithmic game theory*. Cambridge: Cambridge University Press; 2007. Chapter 16. pp. 411–439.
38. Satterthwaite M, Shneyerov A. Dynamic matching, two-sided incomplete information, and participation costs: existence and convergence to perfect competition. *Econometrica* 2007;75(1):155–200.
39. Shen Z-JM, Su X. Customer behavior modeling in revenue management and auctions: a review and new research opportunities. *Prod Oper Manage* 2007;16(6):713–728.
40. Talluri KT, van Ryzin G. *The theory and practice of revenue management*. New York: Springer; 2005.
41. Vulcano G, van Ryzin G, Maglaras C. Optimal dynamic auctions for revenue management. *Manage Sci* 2002;48(11):1388–1407.