Network Market Design Part II: Spectrum Markets

Randall Berry

Abstract

Market-based approaches have promising potential for allocating network resources. Part 1 of this article introduced the game theoretic underpinnings of market design and argued for the need to jointly consider market design with the underlying engineering issues in communication networks. Example research questions in this area were reviewed for wireline networks. In this part, we turn to network market design for wireless systems and in particular for the flexible sharing of wireless spectrum. We use this as a vehicle for discussing challenges that arise in the design of markets in the presence of externalities.

Index Terms

Network Mechanism Design, Game Theory

I. Introduction

Resource sharing is the basic issue underlying both the design of communication protocols as well as the design of economic markets. The first part of this article made the case that these two approaches to resource sharing are not separable and espoused the need for a theory of *network market design*, which combines both economic and technical considerations. Examples of such an approach were discussed in the context of designing markets for bandwidth in a wireline network. In this part, we turn to wireless networks and discuss several other issues that such a theory would need to address. In particular we focus on market-based approaches for *spectrum sharing*.

Currently most licensed spectrum for wireless services is allocated on very coarse scales in both time and space. For example in the U.S, many licenses are allocated for 10 year time-frames and typically cover a large portion of the country. It has been widely argued that this approach has to led to under-utilization of spectrum; this in turn has led to an array of techniques being proposed to enable more efficient use of limited spectrum resources (for example, [1] provides an overview of this area). These techniques include various market structures for trading and/or leasing spectrum on finer temporal and spatial scales (see for example [2]). Here, we use such markets as a way to discuss a number of issues related to network market design, which may also be relevant in other settings.

We highlight a few key differences between wireless spectrum markets and the bandwidth markets considered in Part 1. First, in a bandwidth market, to the first order it is fairly clear how to define the set of assets being allocated. For example, on a single link, the assets are units of bandwidth whose sum is no greater than the link's capacity. Of course, there are secondary considerations. For example, as discussed in Part I, there is still an issue that in practice bandwidth is not infinitely divisible and so there is a choice to be made of how this asset is divided up into indivisible bundles. If temporal dynamics are taken into account there is also an issue of deciding on the time-scale at which allocations occur. However, in the context of wireless spectrum, even such a first order approximation is not clear. For example, one approach to allocating a band of spectrum to a group of users is to require that all users transmit using spread spectrum over the entire band, treating interference from others as noise. With such an approach, the allocation of spectrum might correspond to determining the allowed transmission power of each user or each set of users. Another approach is to use frequency division multiplexing (FDM) and allocate exclusive use of frequency bands to different users. These two approaches lead to very different definitions of the set of assets being traded. These are engineering choices, but the resulting choice also affects any market that emerges. Jointly considering such effects is a key dimension of network market design, which clearly requires both engineering and economic insights.

A second basic property of wireless spectrum is that two users utilizing the same frequency band at nearby locations mutually interfere with each other. Hence, an agent's value for a given "spectrum asset" may depend in part on the allocation of other assets to other agents. Such effects are known in economics as *externalities*. When spectrum is allocated on coarse geographic scales, such interference externalities are a relatively minor problem, in part because the boundaries can be drawn through sparsely populated areas and in part because the boundaries comprise a much smaller portion of the total area. However, if spectrum was allocated on a finer spatial scale, these issues become more relevant. It is well known in economics that the presence of such externalities can greatly complicate market design. We begin in the next section, with a more complete discussion of such externalities. We follow this with a discussion of several approaches for dealing with interference externalities in spectrum markets.

II. EXTERNALITIES AND TRAGEDIES

In the model for bandwidth allocation considered in Part 1, a user's utility is simply a function of the amount of bandwidth she is allocated, and in particular does not depend on the bandwidth allocated to other users. In the case of wireless spectrum, this may not be the case. Due to interference, the utility an agent derives from an allocation may decrease if a "nearby" agent increases his allocation. Economists refer to such effects as *externalities*. This interference effect is a *negative externality*, since increasing the allocation to one agent has a negative effect on the performance of all other agents. Networking problems may also exhibit *positive externalities* when the actions of one agent lead to higher utility of another. An example of this is a peer-to-peer file-sharing system in which users bring resources that help other users as well as themselves. Here we focus on negative externalities.

More formally, as in Part 1, consider a general resource allocation problem of selecting an outcome $x = (x_1, \ldots, x_N)$ for N agents from a set of feasible outcomes X, where x_i represents agent i's share of the resource. Loosely, without externalities, each agent i's utility for a given allocation will depend only on x_i, i.e., it will be a function $u_i(x_i)$ and not depend on the values of x_j, for $j \neq i$. On the other hand, when externalities are present, agents will have utilities that depend on both x_i and x_j for $j \neq i$, i.e., these will be functions of the form $u_i(x_1, \ldots, x_N)$. This is not a precise definition. In particular, note that the choice of labels for each outcome is arbitrary and by simply changing the labels, one can change the dependence of the utilities on these labels. As an extreme example, for any problem, we can simply label each outcome with the corresponding utility received by each agent and so we would have $u_i(x_i) = x_i$ and thus would never see any dependence on x_j in user i's utility.

A more precise definition of externalities is somewhat subtle. To see this, note that even in the case of allocating the bandwidth of a single link, one agent's bandwidth allocation reduces the amount of bandwidth available and so will have an effect on the utility that can be received by other agents. Is this an externality? One approach to answer this is to define an externality not simply in terms of the underlying resource, but in terms of a market for that resource. Given a market, an externality is defined as an effect that one agent causes others, *that is not accounted for in the market* [3]. For example, there are no externalities in a market based on the Kelly mechanism discussed in Part 1 for allocating the bandwidth of a single link, since the dependency of the users via the common capacity constraint is represented via the per-unit price, μ. Though this definition is more precise, the term externality is more often used in the former sense to simply reflect the dependence of one agent's utility function on the resources obtained by another, under some "natural" parameterization of the resources.

To further illustrate the role of a market in determining an externality, consider an "open" market for sharing a single link with capacity C among N agents. In this market, each agent simply requests some amount of bandwidth and is allocated their request, without paying any charge, provided the sum of their requests is less than C; otherwise, they receive nothing. As in Part I, each agent i, receives a utility $u_i(x_i)$ that depends on the amount of bandwidth x_i that she is allocated. Now additionally assume that the agents receive a *disutility* per unit flow, that is proportional to the total traffic on the link, divided by C, i.e., $\frac{1}{C}\sum_i x_i$, which models some form of congestion on the link. Agent i's utility is now given by $u_i(x_i) - \frac{1}{C}x_i(\sum_j x_j)$, where the second term represent a negative externality that is not accounted for in this market. To see the effect of this, further assume that $u_i(x_i) = x_i$ for each agent i. It can be shown that the game defined by this market structure has a unique Nash equilibrium in which each agent requests $\frac{C}{1+N}$ units of bandwidth. For comparison, the efficient allocation in this setting is to give each agent $x_i = \frac{C}{2N}$. The difference between these two allocations is because the market does not provide a way account for the negative costs an agent's allocation has on other agents' utilities, leading agents to over-consume. To see the effect of this, note that the total welfare obtain by the equilibrium allocation is $\frac{NC}{(1+N)^2}$, which goes to zero as N increases. With the efficient allocation, the total welfare is $C/4$, independent of the number of agents, i.e., this market results in an efficiency loss that approaches 100% as the size of the market grows, a situation sometimes referred to as a *tragedy of the commons*.

III. PRICES

One classical "solution" to the previous tragedy is to use a price to "internalize the externality ". In this model, the externality that each user contributes to the overall welfare is given by $\frac{1}{C}x_i\sum_{j\neq i} x_j$. Note that when every other user has the efficient allocation, then the marginal change in this term for each user i is given by $p = \frac{1}{2} - \frac{1}{2N}$. It follows that if we instead charge each user a price of $p = \frac{1}{2} - \frac{1}{2N}$ per unit bandwidth, then the resulting market would have a unique equilibrium in which the users receive the efficient allocation. This price is also known as a

Pigovian tax after the economist Alfred Pigou, who proposed such a solution to externality problems [4]. Note that in this problem all agents were charged a common Pigovian tax of p. This sufficed due to the symmetry of the problem. Without this symmetry, the marginal change in the welfare of the other users due to an change in user i's allocation may depend on i, and so N different Pigovian taxes may be required. Actually determining the optimal Pigovian tax required knowledge of the utilities of each user. Of course, the main motivation for using a market to allocate resources is that this information is not known *a priori*.

To illustrate this difficulty in the context of spectrum, consider a model for sharing a band of spectrum among N agents in a given area. For simplicity assume that each agent corresponds to a single transmitter/receiver pair and that spectrum is shared by specifying the power, P_i that each agent uses to transmit over the common band. Interference from other agents is then treated as additional noise. Each agent i receives a utility $u_i(\gamma_i)$ that is an increasing function of their signal-to-interference plus noise ratio (SINR) γ_i, which is given by

$$\gamma_i = \frac{h_{ii}P_i}{\sigma^2 + \sum_{j\neq i} h_{ji}P_j},$$

where h_{ij} denotes the channel gain from transmitter i to receiver j and σ^2 is the noise power. Here, each user generates a negative externality due to the interference. If users ignore this externality and simply choose their own power P_i to maximize their own utility $u_i(\gamma_i)$, then they would all choose to use the maximum power possible, since their own utility is increasing in their own power. This can result in a total welfare that is much smaller than the optimal. What would Pigovian taxes look like here? They can be expressed in terms of *interference prices* as introduced in [5]. The interference price π_j of user j is the marginal decrease in that user's utility due to an increase in the total interference at that user (i.e., an increase in $\sum_{j\neq i} h_{ji}P_j$). Given the set of interference prices, the Pigovian tax charged to user i (per unit transmit power) is given by the sum of the interference prices from each user $j \neq i$ weighted by the cross channel gains h_{ij}. In general each user will have a different interference price, which depends in part on that user's utility and SINR and these prices will be weighted differently for each user, resulting in a different tax for each user i. Determining these taxes appears to require knowledge of all channel gains as well as the utilities of each agent. However, in [5] it is observed that each agent can compute their own interference price given only local knowledge. Based on this observation, a distributed algorithm is developed in which the agents iteratively exchange interference prices and update powers. With certain restrictions on the class of utilities this is shown to converge to the socially optimal power allocation. Like the Network Utility Maximization (NUM) framework described in Part 1, the interference pricing algorithm from [5] uses these prices mainly as a metaphor to develop a distributed algorithm and is not truly modeling economic incentives. In particular, the agents have an incentive to announce inflated interference prices, since this would reduce their interference and not cost them anything.

The theory of mechanism design discussed in Part 1 of this article provides a framework for dealing with such incentive issues. Indeed, the Vickery-Clarke-Groves (VCG) mechanism can be used in settings with externalities to again provide an incentive compatible and efficient allocation. Recall, for a setting without externalities, i.e., one in which each user's utility only depends on his allocation x_i, this mechanism requires agents to essentially submit their utility function $u_i(x_i)$ and the mechanism is then required to solve an optimization problem corresponding to maximizing the total welfare to determine the allocation and N related problems to determine the payments. In a setting with externalities, the required bids are again the utility functions, only now these must be given as a function of the joint allocation to all agents, i.e., agents must specify $u_i(x_1, \ldots, x_n)$. For the power allocation problem described above, this corresponds to reporting not only one's utility $u_i(\gamma_i)$ as a function of the SINR, but also the relevant channel gains involved in determining the SINR as a function of the powers allocated to each agent. In practice, this could be an excessive amount of information to report and may not even be known accurately (e.g., cross channel gains could be difficult to measure). Once again the mechanism is required to solve $N + 1$ optimization problems: one for the allocation and N for the payments. These tend to be more difficult to solve in the presence of externalities. For example in this power allocation problem, these optimization problems may be non-convex due to the interference. We also note that the VCG payment can be viewed as the total ”externality” that an agent imposes on the welfare of all other agents. This can be contrasted with the Pigovian tax, which gives the marginal change in the externality.

Given the aforementioned difficulties with the VCG approach, it is natural to seek a simpler procedure for pricing such interference externalities, such as the Kelly mechanism discussed in Part 1 for bandwidth allocation. What would be a Kelly-like mechanism for this power allocation problem? In the case of bandwidth, the Kelly mechanism is a *market-clearing mechanism*, i.e., it sets a per unit price so that supply (the links capacity) meets demand (represented by the bids). For this wireless model, the ”supply” is less clear. One way to think about this is that each receiver has a supply of interference it is willing to tolerate. This suggests having nodes bid for the supply of interference at each node, and then again set prices as in the Kelly mechanism so that supply meets demand. Such an approach was studied in [6], [7] for pricing the received interference at only a single ”measurement point” with a fixed supply of interference. In such a setting even if user's are ”price taking,” they are still strategically coupled due to the interference. It is shown in [6], [7], the resulting games for both price anticipating and price-taking users have Nash equilibiria. However, in general this equilibrium is not efficient, even for price taking users and the efficiency loss may be unbounded. In this case, not only is this market clearing mechanism not accounting for ”price anticipating users,” but the resulting single price is not accurately reflecting the externalities among the agents. As we have discussed, truly capturing the externalities requires a different interference price at each receiver. Why not simply run a Kelly-like mechanism for the interference at each of these? There are several difficulties with such an approach. First, how much interference a node is willing to supply may depend on how much power it can use (due to the available interference at other nodes). Second, how much interference a node is willing to buy at one node, would depend on how much it may buy at other nodes, since its total power is determined by the smallest ”interference allocation” it receives. Finally, implementing such an approach would require the seller to have knowledge of the received interference at each receiver, which may not be feasible in practice.

Instead of pricing the received interference at each receiver. An alternative market for this scenario is to view the supply as the potential transmission power available at each node i. Each unit of this power can be allocated to either node i, allowing her to increase her transmission power or to another agent j, which prevents i from using this power and thus reduces the interference at node j. A market for such a scenario is considered in [8], in which discrete units of power at one user were allocated to two users sequentially using a second-price auction for each unit. It shown that in the worse case such a mechanism can have an efficiency that goes to zero as the number of discrete units increases. However numerical results indicate that in average cases, there may be little if any efficiency loss. Extending the analysis in [8] to more than 2 users appear to be difficult, in part because whenever one agent $j \neq i$ is allocated a unit of agent i's transmission power, it reduces the interference for all agents other than i, which gives these agents an incentive to "free-ride" on each other.

IV. Asset Design

Another "solution" to the externality problem is to design spectrum assets to limit or remove externalities. To illustrate this, we return to the power allocation problem introduced above. Instead of sharing the band of spectrum by having each user spread their signal over the entire band, we could instead subdivide the band into orthogonal frequency bands (or assign users orthogonal time-slots). With such a change, the spectrum allocation problem becomes equivalent to a bandwidth allocation problem as in Part 1, and so there are no longer any externalities present. This leads to simpler market design, but from an engineering perspective it may not be ideal. In particular, we know that if the transmitter/receiver pairs are far enough apart, then the earlier "power sharing" model will better exploit frequency re-use and result in better spectrum utilization. This illustrates an under-explored trade-off between designing assets so as to have simpler market mechanisms versus a more efficient allocation. As another example of this trade-off, one can obtain an even simpler market design by simply allocating the entire spectrum band to one agent, e.g., via a second price auction, but again if each agent only requires a small amount of the available spectrum, this could be very inefficient.

A practical approach to navigating this trade-off is to instead allow agents to use the same frequency band when they are far enough apart, but require nearby users to use different bands. When users sharing the same band are far enough apart, the externalities will become insignificant and can be safely ignored. For example, suppose a spectrum asset is defined as the right to transmit in a given area with a fixed power mask in a given frequency-band. We can then represent a set of C such assets as nodes in an interference graph, $\mathcal{G}$, so that two assets are connected by an edge if they significantly interfere with each other (see Figure 1). If we wish to allocate assets without significant interference, then this corresponds to only allocating assets that belong to an independent set in $\mathcal{G}$ (here we focus on simply allocating a single frequency band, with multiple bands a different independent set could be used for different frequency bands). If we simply fix one independent set for all time and only allocate these assets, then the assets would have no externalities. However, the "best" choice of assets, i.e., the independent set with the most value may differ from time-to-time depending on the agent's demands. An alternative approach would be to dynamically determine the independent set via a market mechanism. Given the true value for each asset, the efficient allocation is to find a maximum weight independent set in $\mathcal{G}$, where the weights correspond to the values. Unfortunately, this is an NP-hard problem, and thus implementing such an allocation via a VCG mechanism may be computationally

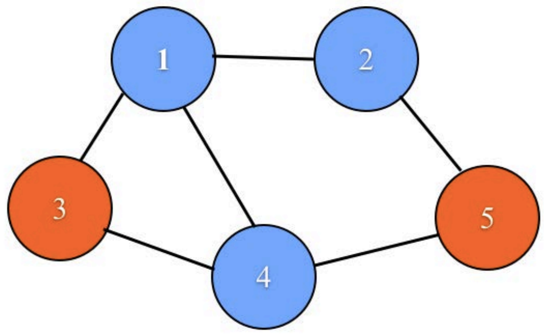

Fig. 1. Example of an interference graph $\mathcal{G}$ for 5 spectrum assets. One independent set is shown in orange.

challenging. Of course, the importance of such computational considerations depend on the size of problem and the time-scale at which allocations are performed. If this complexity is a concern, then instead of exactly solving for the maximum independent set, one could instead employ an approximation algorithm; however, in such cases the incentive properties of the VCG mechanism may no longer hold.

Designing approximation schemes for general combinatoric auctions that preserve the VCG's incentive properties is a topic that has received much attention in the algorithmic game theory literature (e.g., see [9]). In the context of spectrum markets, such question have also received some attention, e.g. [10] develops a design called VERITAS. VERITAS applies to the preceding setting in the special case where each bidder wants a single asset from C (more generally it applies to settings with multiple frequency bands). Given an agents value for its asset, VERITAS finds a independent set in time less than $O(n^3)$. Moreover this mechanism is *truthful*; i.e., agents have no incentive to not bid their true value. The key to showing this is establishing that the allocation rule is *monotonic*, meaning that if bidder i wins an assets when bidding b_i, it will continue winning if it raises its bid to any value greater than b_i, assuming all other agents keep their bids fixed. For setting such as this in which each agents desires a single asset, such monotonicity is known to be necessary and sufficient for a mechanism to be truthful [11]. The cost for the computational tractability of this mechanism is that it is no longer be guaranteed to find the optimal allocation (if it did, it would be finding a maximum weight independent set in polynomial time). No approximation bound for VERITAS is given in [10], but the mechanism is shown to perform well in numerical results. Another example of a truthful, computationally tractable mechanism is given in [12], which also allows for agents to share the available assets. Again in this work agents value a single asset (or share of an asset) and so monotonicity suffices to show truthfulness. In the more general setting where agents may desire multiple spectrum assets, it becomes more difficult to characterize monotonicity and to design truthful tractable mechanisms. Relatively little has been done in this space in the context of spectrum markets.

When representing spectrum asset in terms of an interference graph, a fundamental underlying question is how to determine the level of interference required for link to be present in this graph. This level could vary greatly with the intended application and technology used by a bidding agent, making it difficult to predetermine. An alternative to this, introduced in [13] is to not allocate only independent sets but instead allow every set of asset to be allocated. In such a setting, suppose that an agent wanted an asset i but only if there was no interference from a neighboring asset j. The agent could achieve this by simply purchasing both assets i and j and simply not using j. On the other hand another agent who wanted asset i but could tolerate interference from j would only need to purchase asset i. With such an approach, spectrum assets have *complementarities* meaning that the value of a bundle of assets may be greater than the sum of the value of the individual assets. Such an approach will be more efficient than restricting allocations to be independent sets. A model for such a setting is studied in [13], for which it is shown that the underlying allocation problems are still NP-hard, unless the interference graph is restricted to be a line or ring. Little work appears to have been done to develop incentive compatible approximation algorithms for such a setting.

In the previous section we considered allocation spectrum in term of assigning powers to transmitting nodes, while in this section we focused on models in which power is fixed and one agents has the right to use this in a given spatial region. We briefly mention an approach from [13] that combines these two ideas. Namely, a market is considered in which agents bid for both a spatial region and the power emitted from that region, which effectively determines the boundary of the region. For a simplified model, it is shown that adding this flexibility actually make the resulting allocation problem computationally tractable. More precisely, this model applies to a setting in which an agent's value for an asset is proportional to the "radius" of the spatial region served. Under this assumption, the optimal boundaries for a given assignment can be found by solving a linear program. Other ways of defining spectrum assets are also possible. For example the definition could require nodes to follow a protocol like CSMA, which results in a different form of interference externality. Ideas of "cognitive sensing" could also be incorporated into the definition, perhaps leading to both primary and secondary spectrum assets. Finally, more advanced physical layer techniques like the use of multiple antenna could play a role. However, defining a spectrum asset in a way that is tied to a specific technology also has a risk, namely, it may make it difficult for alternative competing technologies to emerge. A clear picture of the costs and benefits of such approaches are yet to emerge.

So far we have focused on a market for trading "raw spectrum," i.e., access to the physical medium. Another question related to asset design is whether this is the right layer at which to run a dynamic market. If a new entrant wishes to offer a wireless service, it is not clear that it is best served by acquiring raw spectrum and deploying its own infrastructure. It could instead contract with existing "infrastructure providers" to offer its service over their infrastructure. Such infrastructure providers could have a relatively static supply of raw spectrum on which they build a flexible infrastructure, which is used to provide transport services to various higher layer agents. The resulting market would then be for these transport services and the market design issues would be closer to those for wired bandwidth discussed in part 1. To a limited extent, this type of market exists today for example when cellular service providers sell access to their networks for various Mobile Virtual Network Operators (MVNOs). Enabling such a market on a much more extensive scale could help facilitate more efficient spectrum usage with requiring trading of raw spectrum. An advantage of this approach is that if infrastructure providers own spectrum at neighboring locations, they can manage the interference externalities themselves, making the market design issues simpler. Potential disadvantages are that such a market structure might not provide as much opportunity for new physical layer services to emerge and the transport services offered by the infrastructure providers may not be well suited for all applications. Again, more work is needed to better understand and model such trade-offs.

V. Bargaining

The final approach for dealing with externalities that we will discuss is via *bargaining*. To motivate this consider allocating power to two neighboring bands of spectrum owned by two different wireless service providers. The two providers interfere with each other creating an externality. We can model this by assuming that provider i's utility depends on both the power he is allocated and the power allocated to the other provider. Under a bargaining approach, instead of attempting to price this externality or design assets in such a way that the externality does not arise, we simply allocate the bands to both providers and let them bargain with each other to determine the needed power allocation. For example, suppose that provider 1 sees a gain of v_1 from provider 2 reducing his power and provider 2 only loses $v_2 < v_1$ from doing this. Then, if provider 1 offers provider 2 a payment of p, where $v_2 < p < v_1$ to do this, provider 2 should be willing to accept this and the overall welfare will improve. This is an example of what is known as the *Coase theorem*, which states that if trade for an externality can occur, then bargaining will lead to an efficient outcome, independent of the initial allocation. This result is attributed to the economist Ronald Coase, who in fact developed it while studying the allocation of wireless spectrum in the 1950's [14].

A prerequisite for the Coase theorem to apply is the existence of well-defined and enforceable *property rights* for the assets being traded. In other words, agents must be able to clearly value the assets and ensure that the counter party follows through on their side of the bargain. In the context of spectrum ensuring well defined property rights may be subtle, as illustrated by the recent dispute over LightSquared's plans to offer nationwide 4G LTE service. At the heart of this dispute is the FCC requirement that services deployed in one band do not cause ”harmful interference” to services in adjacent bands. This does not provide a well-defined property right, as harmful interference depends in part on the type of receivers deployed in the adjacent bands. The key point here is that the definition of a property right has to clearly define the underlying externality in order for agents to able to bargain over it. In this spirit, there has been a line of work by the legal and policy communities on defining property rights to enable spectrum markets (e.g. [15]). Once again this is an issue that is not cleanly separable from technical considerations, though the technical community has provided little input.

The Coase theorem may seem to suggests a very simple market structure for well defined assets; simply allocate the assets arbitrarily with well defined property rights and let the agents bargain. The problem with this conclusion is that in practice there are a number of *frictions*, which may impede efficient bargaining. For example, there may be costs and delays involved in bargaining. These delays may be due to the time it takes to agree on a bargain as well as the search time needed to find a counter party to bargain with. Moreover these tend to increase with the number of parties that one needs to bargain with. In the case of spectrum markets, this suggests that bargaining may be more difficult when spectrum is allocated on finer temporal and spatial scales; finer temporal scales restrict the time for bargaining and finer spatial scales may increase the number of relevant counter parties. Another friction in bargaining is due to *imperfect information*, i.e., when agents do not exactly know each other's valuations. Returning to the previous example, suppose that $v_1 = 3$ and $v_2 = 2$ so that provider 1 should be willing pay some price p slightly more than 2, which will result in provider 2 lowering his power and the overall welfare improving. Now suppose that provider 1 only knows that v_2 is distributed uniformly on the interval $[0, 3]$; in this case provider 1's expected benefit is maximized if she makes provider 2 an offer of 1.5, which would not be accepted, even though the overall welfare would improve if trade occurred. Of course, this is just one possible way the two agents can bargain; however, a result due to the economists Myerson and Satterthwaite [16] shows that the underlying problem is fundamental: namely in the presence of incomplete information there is no bargaining procedure that can always guarantee the efficient outcome. This discussion suggests that simply relying on bargaining alone may not be sufficient; markets can be used to reduce these frictions and provide incentives for agents to reveal private information. Of course, even with a well designed market, there may still be opportunities for agents to bargain after receiving their allocations. Incorporating such considerations in the design of dynamic spectrum markets is another area that has not been fully developed.

VI. Conclusions

From both parts of this paper it should be clear that network market design contains a number of challenging and under-explored topics. In this part we have focused on a number of such issues for spectrum sharing markets that must operate in the presence of interference externalities. In this part, we have focused mainly on single-sided markets in which one "spectrum manager" is selling access to a set of service providers. As discussed in Part 1, market design becomes more complicated for doubled-sided exchanges in which there are both multiple buyers and sellers. Such a setting could also arise in spectrum markets, e.g., when multiple license holders seek to lease their spectrum to multiple secondary service providers.

As in Part 1, we have largely ignored temporal dynamics, which raises a number of additional questions. For example, the price of spectrum assets will change based on time-varying demand. More elastic users can attempt to exploit this by deferring their usage to cheaper times; however their ability to do this will be based in part on how well they can predict future usage, which is determined by the actions of other strategic agents. Less elastic users may want to guarantee some level of service in the face of uncertain future price fluctuations. Perhaps, a "spectrum futures" market could be used to meet this demand. Analyzing such settings is challenging, in part because the underlying economic theory for dynamic environments is much less developed.

References

[1] Q. Zhao and B. Sadler, "A Survey of Dynamic Spectrum Access," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 79-89, May 2007.

[2] R. Berry, M. Honig and R. Vohra, "Spectrum markets: motivation, challenges and implications," *IEEE Communications Magazine,* vol. 48, no. 11, pp. 146-155, Nov. 2010.

[3] A. Mas-Colell, M. D. Whinston and J. R. Green, *Microeconomic Theory*, Oxford University Press, 1995.

[4] A. Pigou, *The Economics of Welfare*, Macmillan, London 1920.

[5] J. Huang, R. Berry and M. L. Honig, "Distributed Interference Compensation for Wireless Networks," *IEEE on Journal on Selected Areas in Communications,*, vol. 24, no. 5, pp. 1074-1085, May 2006.

[6] J. Huang, R. Berry, and M. Honig, "Auction-based Spectrum Sharing," *Journal of Mobile Networks and Applications (MONET)*, vol. 11, no. 3, pp. 405-418, June 2006.

[7] R. Berry, "On proportional power sharing mechanisms for secondary spectrum markets," in *Proc. of 3rd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pp. 13-16, Dec. 2009.

[8] J. Bae, E. Beigman, R. Berry, M. Honig, and R. Vohra, "Sequential Bandwidth and Power Auctions for Distributed Spectrum Sharing," *IEEE Journal on Selected Areas in Communication*, vol. 26, no. 7, pp. 1193-1203, September 2008.

[9] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, *Algorithmic Game Theory,* Cambridge University Press, 2007.

[10] X. Zhou, S. Gandi, S. Suri, and H. Zheng, "eBay in the Sky: Strategy-proof wireless spectrum auctions," in *Proc. of the 14th ACM International Conference on Mobile Computing and Networking*, pp. 2-13, 2008.

[11] R. Myerson, "Optimal Auction Design," *Mathematics of Operations Research*, Vol. 6, pp. 58-73, 1981.

[12] I. Kash, R. Murty, and D. Parkes, "Enabling Spectrum Sharing in Secondary Market Auctions," in *6th Workshop on the Economics of Networks, Systems, and Computation (NetEcon)*, 2011.

[13] H. Zhou, R. Berry, M. Honig and R. Vohra, "Spectrum Markets with Interference complementarities," in *Proc. of 9th Intl. Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks(WiOpt)*, Princeton, NJ. May 9-13, 2011.

[14] R. Coase, "The Problem of Social Cost," *Journal of Law and Economics*, Vol. 3, pp. 1-44, Oct. 1960.

[15] G. Faulhaber and D. Farber, "Spectrum Management: Property Rights, Markets and Commons," in *Rethinking rights and regulations: Institutional responses to new communication technologies*, MIT press, Cambridge MA, 2003.

[16] R. Myerson and M. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, vol. 28, pp.265-281, 1983.