

Spectrum Pooling with Competitive Service Providers

Melina Muthuswamy[†], Randall Berry[‡], Michael Honig[‡], Thanh Nguyen*, Vijay Subramanian[§],
Rakesh Vohra[†]

[†]Univ. of Pennsylvania, [‡]Northwestern Univ., [§]Univ. of Michigan, *Purdue

Abstract—We consider a model in which two competing wireless service providers with licensed spectrum may pool a portion of their spectrum to better exploit statistical multiplexing. Given an amount of pooled spectrum, the providers engage in Cournot competition. We study the impact of pooling spectrum on the outcome of this competition and show that the gains from multiplexing are dissipated due to the competition among the providers.

I. INTRODUCTION

Traditionally, wireless service providers (SPs) have used their own bands of licensed spectrum to offer service to their customers. *Spectrum pooling* is an alternative approach in which the same spectrum is shared by multiple SPs offering the same services [1], [2].¹ Each SP accessing a spectrum pool could do so using its own infrastructure perhaps coordinated via a spectrum manager [4]. This coordination could also occur through using a technology such as Licensed Assisted Access (LAA) [5]. Alternatively, this could be implemented via radio access network (RAN) sharing [6] in which case infrastructure would be shared along with spectrum.

Spectrum is inherently a congestible resource in which the quality of service that can be provided to a user in a given band declines as more users utilize that band. A key advantage of pooling spectrum is that it can decrease congestion in the presence of randomly varying demands. In particular, pooling spectrum enables SPs to exploit statistical multiplexing similar to the effects that arise in traditional wired networks [7]. However, in a competitive market, the benefits of improved statistical multiplexing will influence the competition among the SPs who pool their spectrum. Our goal in this paper is to understand the impact of such competition on the gains from pooling spectrum.

We consider a scenario in which there are two competing SPs, each having its own licensed spectrum. With randomly varying and independent demands presented to the two providers, congestion is minimized by pooling the two bands which effectively *increases* the joint capacity of the SPs by enabling them to exploit statistical multiplexing.

A reduction in congestion will have an impact on price and demand. As the reduction in congestion is akin to an increase

in capacity, we would expect lower prices. This will increase both demand and congestion. Will the increased demand offset the drop in prices enough to make the providers better off? While more consumers are being served, they may face greater congestion – are they better off?

We study these pooling effects on SP revenue and consumer surplus in a static scenario where the cost of congestion in a particular band is given by a weighted combination of mean load (customers served normalized by bandwidth) and its variance, where the latter term captures the benefits of statistical multiplexing. That is, for a given mean load, the congestion cost decreases with bandwidth due to the decrease in normalized variance.² To allow the possibility of partial pooling while maintaining competition, we consider a scenario in which the two providers agree to share a fraction of their allocated spectrum, and then split the users across their shared and proprietary bands to equalize the corresponding congestion costs.

Competition between the SPs is captured via a Cournot model: the SPs simultaneously choose the load they wish to serve and the total load determines the market price. The higher the total load, the lower the market price. Consumers make the decision to purchase service based on the *delivered price* which is the market price plus congestion cost. This model extends the model of competition in [8] by incorporating statistical multiplexing gains. That paper also examines the benefits of sharing spectrum. However, the benefit of sharing in that paper does not arise from the SPs pooling some of their resources but rather from sharing intermittently available spectrum from an outside source.

We find that the gains from multiplexing are dissipated in the competition between the SPs. In fact the SPs profits may be lower with pooling which eliminates their incentives to share. More consumers are served because prices are lower but congestion is higher than before. Consumer surplus, however, increases with the amount of bandwidth that is shared. We also compare this with a scenario in which there is a single monopoly SP with two bands of spectrum. In the monopoly case, pooling the two bands leads to both higher profits and greater consumer surplus compared to when the monopolist keeps the two bands separate.

R. Berry was supported in part by NSF grants CNS-1908807 and ECCS-2030251. V. Subramanian was supported in part by NSF grants ECCS-2038416, CCF 2008130, and CNS 1955777.

¹This type of approach is also called *co-primary spectrum sharing* [3].

²The reduced variability is best ascertained by looking at higher moments. Here we restrict to just the variance but remark that our choice captures the first two terms of the effective bandwidth formula in [7].

A. Prior Literature

This paper is a contribution to the literature on competition among wireless service providers with shared spectrum (e.g. [9]–[12]). In particular, it builds on a line of literature that adapts models for competition among firms with congestible resources (e.g. [13]). That literature has not considered the impact of competition on pooling to exploit multiplexing gains.

There is an extensive literature on technologies for pooling spectrum to harvest a multiplexing gain (e.g. [2], [14], [15]). In this paper we take the existence of multiplexing gains as given and capture the underlying technology via a congestion function. Our focus is on whether SPs have an economic incentive to exploit such gains by pooling their resources.

Another motivation for pooling spectrum is that it may reduce the cost of deploying infrastructure (see e.g. [16], [17]). Here, we assume that any investment in infrastructure is fixed *a priori* and focus solely on the multiplexing benefits of pooling given the fixed investment.

There has also been a line of work on the design of incentives for SPs to enter into pooling agreements (see e.g. [17], [18]). We take such an agreement as given here and do not explicitly model this. However, by comparing the profits of SPs when they pool spectrum, our results provide insights into when such agreements may emerge.

II. MEAN-VARIANCE MODEL OF LATENCY

To capture statistical multiplexing, we consider a mean-variance model for the latency of users within a given band of spectrum. Suppose a single band of capacity C serves a mass x of non-atomic users, where the capacity C is proportional to the bandwidth of the spectrum.³ We model the latency per user as

$$\ell(x; C) = \frac{x}{C} + \frac{xv}{C^2}, \quad (1)$$

where v is a parameter related to the variability of traffic on the band. We can think of this expression as arising from a setting where each user brings a random quantity of traffic at a given time with a unit mean and a variance of v , which is independent across the users. The system load (given by the amount of traffic divided by the capacity, C) will then have a mean of x/C and a variance of xv/C^2 . Our latency model assumes that users are sensitive to the sum of the mean and the variance of the system load.⁴

We now show that this model of latency provides gains from multiplexing. If we share by combining two bands of capacity C_1 and C_2 which serve x_1 and x_2 customers, respectively, the latency on the combined band is:⁵

$$\ell(x_1 + x_2; C_1 + C_2) = \frac{(x_1 + x_2)}{(C_1 + C_2)} + \frac{(x_1 + x_2)v}{(C_1 + C_2)^2}. \quad (2)$$

³The capacity can also depend on the technology deployed in the band, which here we assume is fixed across all bands.

⁴More generally this could be a weighted sum where the weight for the variance term can simply be absorbed into v .

⁵Depending on the approach used to pool these two bands, the combined capacity could be less than $C_1 + C_2$, e.g., there may be a degradation due to coordinating different SPs use of the spectrum. We ignore such effects here.

Suppose $x_1 = x_2 = x$ and $C_1 = C_2 = C$. Then, we have

$$\frac{x}{C} + \frac{xv}{C^2} > \frac{2x}{2C} + \frac{2xv}{(2C)^2} = \frac{x}{C} + \frac{xv}{2C^2}.$$

Here, the term on the left is the latency experienced if two separate bands of capacity C are each used to carry x units of traffic. The term on the right is the latency if the two bands are pooled creating a band with capacity $2C$, which is used to serve the combined $2x$ units of traffic. The inequality establishes that there is a multiplexing gain, which is increasing with the variance in the traffic v . As v grows, this gain approaches 50%. When $v = 0$ (the latency model used in [8]), there is no multiplexing gain.

III. MONOPOLY

As a baseline we consider a monopoly SP under the mean-variance latency model with a total capacity of $2C$. We assume the SP can operate on two *separate* bands having capacity C_1 and C_2 where $C_1 + C_2 = 2C$ or on a single band with capacity $2C$. We compare the revenue, latency, and price across the two scenarios.⁶ For the purposes of comparison of different market structures, in our analysis below we will modify the mean term of the latency to be $\beta x/C$ for parameter $\beta \in \{0, 1\}$.⁷

We assume a pool of infinitesimally small customers each of whom bears a latency cost as in (7) that depends on the traffic carried on the band which serves them. Customers care about the delivered price, p_d , which is the posted price plus the latency cost incurred and related to demand by the inverse demand curve of $p_d = 1 - x$ where $x \in [0, 1]$ is the normalized mass of customers accepting service.⁸ Thus, the SP's revenue per customer will be the difference between p_d and the latency suffered.

The surplus obtained by the y th user receiving service is given by the difference between a consumer's value for service (given by the inverse demand curve evaluated at y) and the delivered price that consumer experiences (p_d). The overall consumer surplus is given by integrating this over all users receiving service. This results in a surplus of $(1/2)(1 - p_d)^2$. From this it is clear that a lower delivered price results in a larger consumer surplus.

Pooling allows the monopolist to exploit multiplexing which means they can carry the same volume at lower latency. In turn this means the monopolist can raise their posted price leaving the delivered price unchanged. Alternatively, they can also serve a larger volume of customers at the same announced price.

A. Separate Bands

We first consider the mean-variance monopoly model with two separate equal-sized bands. If x is the amount of traffic carried on each band, then, the total traffic is $2x$ and delivered

⁶We can also interpret this as two distinct SPs that coordinate on the traffic they carry both before pooling their separate bands as well as after.

⁷Nominally, we will set $\beta = 1$, but will consider the case of $\beta = 0$ when comparing to the variance-only results in Section IV-B.

⁸Note, in our model price influences the number of customers accepting service but not the amount of traffic a customer generates.

price must be $1 - 2x$. Latency per customer will be $\beta \frac{x}{C} + \frac{xv}{C^2}$. Hence, total profit will be

$$\pi = (2x)(1 - 2x) - \beta \frac{2x^2}{C} - \frac{2x^2v}{C^2}. \quad (3)$$

The profit is maximized by setting $x = \frac{C^2}{4C^2 + 2\beta C + 2v}$, giving a maximum profit of

$$\pi = \frac{C^2}{4C^2 + 2\beta C + 2v},$$

which interestingly is equal to $1/2$ of the optimal quantity served (given by $2x$). The profit maximizing *delivered* price and the corresponding *total* latency are given by:

$$p_d = 1 - 2x = \frac{C^2 + \beta C + v}{2C^2 + \beta C + v},$$

$$l = \frac{C^2 (\beta C + v)}{2(2C^2 + \beta C + v)^2}.$$

Per user latency will be

$$\frac{\beta C + v}{4C^2 + 2\beta C + 2v}.$$

Subtracting the per user latency from the delivered price p_d gives that the market price is $1/2$, which is the same as what a monopolist would charge if faced with the same demand curve in a market without latency. Here, latency causes the monopolist to reduce the quantity of users it serves but to do so in a way that keeps the market price fixed.

B. Pooled Bands

Now suppose the two separate bands of capacity C are pooled into a single band of capacity $2C$. If x is the total traffic carried on the pooled band, the maximum profit of the SP will be given by

$$\max_x x(1 - x) - \beta \frac{x^2}{2C} - \frac{x^2v}{4C^2}. \quad (4)$$

The optimal value of x is $\frac{2C^2}{4C^2 + 2\beta C + v}$ resulting in a profit of

$$\pi = \frac{C^2}{4C^2 + 2\beta C + v},$$

which again is equal to $1/2$ of the total traffic served. The profit maximizing delivered price and the corresponding total latency are shown next.

$$p_d = 1 - x = \frac{2C^2 + 2\beta C + v}{4C^2 + 2\beta C + v},$$

$$l = \frac{C^2 (2\beta C + v)}{(4C^2 + 2\beta C + v)^2}.$$

Per user latency will be

$$\frac{2\beta C + v}{2(4C^2 + 2\beta C + v)}.$$

Notice that the per user latency has declined compared to the case of separate bands.

Once again the market price is given by subtracting the per user latency from the delivered price, which gives $1/2$. Hence, with pooling, a monopolist charges the same market price, but the reduced latency due to pooling enables it to expand the number of users served, which in turn increases its profit. Further, since more customers are served at the same market price, both the per user latency and the delivered price decrease. This means that consumer surplus will also increase, so that both the monopolist and consumers benefit from pooling.

Note also that if $v = 0$, the outcome in the pooled and separate cases become the same, reflecting the fact that when $v = 0$ there is no multiplexing gain. As v grows it can be seen that the ratio of the profit with pooling to that without pooling also grows and approaches a factor of 2, i.e., for large values of v a monopolist can nearly double its profit through pooling.

IV. COURNOT MODEL

We next turn to a model of two SPs competing to provide service to a common pool of users. Each SP has its own proprietary band of spectrum and can contribute a portion of its band for shared use by both SPs. Each SP can then serve customers on both its proprietary band and the shared band. We again assume a pool of infinitesimally small customers with a downward sloping *linear* inverse demand curve. Each customer served bears a latency cost that depends on the traffic carried by a SP on the both the proprietary and shared band. SPs compete a'la Cournot by choosing the quantity of users to serve.

Suppose C_i is the capacity of firm i 's proprietary band ($i \in \{1, 2\}$) and let $\frac{\alpha}{2}$ be the amount of bandwidth each firm contributes for shared use. Notice, we assume that each firm contributes the same amount. The total capacity of the shared band is then α .

Denote by $x_{p,i}$ the amount of traffic SP i serves on its proprietary band and let $x_{s,i}$ be the traffic it serves on the shared band. The total congestion cost to SP i (summed across all users) is then

$$\ell_i = x_{p,i} \ell_{p,i} + x_{s,i} \ell_{s,i} \quad (5)$$

where

$$\ell_{p,i} = \frac{x_{p,i}}{\left(C_i - \frac{\alpha}{2}\right)} + \frac{x_{p,i}v}{\left(C_i - \frac{\alpha}{2}\right)^2}, \quad (6)$$

$$\ell_{s,i} = \frac{\sum_{j=1}^2 x_{s,j}}{\alpha} + \frac{v \sum_{j=1}^2 x_{s,j}}{\alpha^2}. \quad (7)$$

Since the latency function in both the proprietary and shared bands has the same functional form in terms of quantity and capacity, the underlying assumption is that the SPs deploy the same technology in all bands. One can interpret this form of the latency function as arising from SP i randomly assigning each unit of traffic to its proprietary band or to the shared band. Specifically, the probability that SP i assigns a customer to its

proprietary band is $f_{p,i} = x_{p,i}/(x_{p,i} + x_{s,i})$. The total latency cost on this band is then $f_{p,i}\ell_{p,i}(x_{p,i} + x_{s,i})$, and likewise for the shared band, which is consistent with (5).

The actual price paid by customers is the difference between the delivered price and per user latency. The delivered price, p_d , given the pair $(x_{p,i}, x_{s,i})$ for each SP i is

$$p_d = 1 - \sum_{i=1}^2 (x_{p,i} + x_{s,i}).$$

Therefore, SP i 's profit will be

$$\pi_i = \max_{x_{p,i}, x_{s,i}} (x_{p,i} + x_{s,i}) \left(1 - \sum_{j=1}^2 (x_{p,j} + x_{s,j}) \right) - \ell_i. \quad (8)$$

Note that the profits of the two SPs are coupled through both the inverse demand term and through the latency on the shared band. Hence, we consider the (Nash) equilibrium of the resulting game in which each SP i selects its quantities $(x_{p,i}, x_{s,i})$ to maximize its own profit as in (8). Fixing the quantities of the other provider, the objective in (8) is concave in $(x_{p,i}, x_{s,i})$ and so we can characterize each SPs best responses using first order optimality conditions. This results in a set of linear equations whose solution is the market equilibrium. In the following, we will present numerical plots of the resulting equilibrium quantities and omit the detailed derivations due to space considerations.

When $\alpha = 0$, the capacity of the shared band is 0. Since neither firm can place any traffic on this band, the terms associated with α in (7), which enter (8), vanish and $x_{s,i} = 0$ for all i . We will be interested in how equilibrium profits, prices, and consumer surplus change as α increases from 0 up to $2 \min\{C_1, C_2\}$. Figures 1-4 illustrate the equilibrium behavior as α increases for the case $C_1 = C_2 = 1$ and $v = 0.5$. The largest feasible value of α is 2. At this value, no SP has any proprietary bandwidth. For comparison, we also show the monopoly outcomes from Section III as two red asterisks in each figure. The mark at $\alpha = 0$ corresponds to the case where the monopolists keeps the two bands separate, and the mark at $\alpha = 2$ corresponds to the case where the monopolist pools the two bands.

For this setting, the equilibria are symmetric, i.e., the two SPs select $x_{p,1} = x_{p,2}$ and $x_{s,1} = x_{s,2}$ so that they are both putting the same amount of traffic on the shared bands and the same amount on the proprietary bands.⁹ Furthermore, as the size of the shared band increases, the quantities served on the proprietary bands decrease and the quantity served on the shared band increases.

Figure 1 illustrates how the delivered price changes with α . It rises initially and then declines. This means that the volume of users initially falls and then rises. Comparing the delivered price at $\alpha = 0$ and $\alpha = 2$ with that of a monopolist, we

⁹More generally, if the two providers have the same capacity, the equilibrium will be symmetric. If their capacities differ they will place different amounts of traffic on their proprietary bands, but the same amount on the shared band.




Fig. 1: Delivered Price(\$) vs α when $C_1 = C_2 = 1$ and $v = 0.5$.




Fig. 2: Profit(\$) vs. α when $C_1 = C_2 = 1$ and $v = 0.5$.

see that in both the competitive and monopolistic cases, the delivered price is lower when $\alpha = 2$. At both values of α , competition results in a lower delivered price.

Figure 2 shows that SP profit is essentially flat as α increases. In the monopolist case, the profits slightly improves for $\alpha = 2$ compared to $\alpha = 0$, consistent with the discussion in Sect. III.¹⁰ Consumer surplus displays an increase as shown in Figure 3. In the monopoly case, there is also a gain in consumer surplus from pooling, but this is much smaller than the gain in the competitive case. Figure 4 shows that total latency (summed across users) initially declines and subsequently increases.

These figures suggest that if there is a multiplexing benefit for SPs it may be slight and only for α close to zero. For moderate or even large values of α , only users benefit because consumer surplus is increasing. Note that latency and consumer surplus both increase. This means that the benefits to users come in the way of lower prices rather than

¹⁰Here the value of v is relatively small and so the gain from pooling for the monopolist is also small.




Fig. 3: Consumer Surplus(\$) vs. α when $C_1 = C_2 = 1$ and $v = 0.5$.




Fig. 5: Delivered Price(\$) vs. α for mean-only model.




Fig. 4: Latency vs. α when $C_1 = C_2 = 1$ and $v = 0.5$.

lower latency. This is because multiplexing, by introducing additional capacity, gives the firms an incentive to increase the volume of traffic they serve. This results in a lower delivered price. While more consumers are served at a lower delivered price, their latency increases. Thus, SPs don't benefit from multiplexing. This is seen most clearly when we compare the outcome here to the monopoly case.

To verify this intuition, we next examine two polar cases of the model. The first is when $v = 0$ (mean latency only), i.e., there are no gains from multiplexing. The second is where latency depends on the v term only (variance only), which represents the case with the largest gains from multiplexing.

A. Mean-Only Cournot Model




Here, we consider the mean-only variation of the Cournot model as a no multiplexing gain baseline. We model this by removing the variance terms in Equation (8) by setting $v = 0$. The profit optimization problem of firm i is:

$$\max_{x_{p,i}, x_{s,i}} (x_{p,i} + x_{s,i}) \left(1 - \sum_{j=1}^2 (x_{p,j} + x_{s,j}) \right) - \frac{x_{p,i}^2}{C_1 - \frac{\alpha}{2}} - x_{s,i} \left(\frac{\sum_{j=1}^2 x_{s,j}}{\alpha} \right). \quad (9)$$

Again, $x_{p,i}$ denotes the quantity placed on the proprietary band of SP i , $x_{s,i}$ is the quantity placed on the shared band by SP i .

Figures 5-8 depict how the equilibrium changes as α increases for the case $C_1 = C_2 = 1$. As with the full mean-variance model the largest feasible value of α is 2 and at this value neither SP has any proprietary bandwidth. The monopoly outcomes are again indicated via red asterisks. In the mean only model, the monopoly outcome is the same for both the separate and the pooled cases as there is no multiplexing gain. Hence, in each figure the two monopoly outcomes shown are identical.

Again, each SP will serve the same number of customers on their individual proprietary bands and the same number on the shared band. Figure 5 illustrates how the delivered price changes with α . It initially falls and then rises. This means that the volume of customers served initially rises. Since the firms are serving more traffic at lower prices, as α increases, the latency also increases causing a rise in prices. This rise in latency is observed in Figure 8. Moreover, the sharing between the firms is what causes this observed rise in latency. Once the latency increases to a certain point, the firms have an incentive to cut back on volume of customers served. Figure 6 shows a limited change in profit and Figure 7 shows a slight rise and fall in consumer surplus. The initial rise in consumer surplus can be attributed to the fact that as α increases more customers are being served. As α increases to near total sharing, consumer surplus declines. Compared to the monopoly case, again profits are lower with competition and consumer surplus is larger.

Fig. 6: Profit(\$) vs. α for mean-only model.Fig. 7: Consumer Surplus(\$) vs. α for mean-only model.Fig. 8: Latency vs. α for mean-only model.

Visually, there is little to distinguish sharing under the mean only model from the mean and variance model, suggesting that competition eliminates any potential benefits of multiplexing.

B. Variance Only Model

Now, we examine the variance-only variation of the model in order to see the effects of multiplexing unencumbered by the mean term. The profit of firm i will be

$$\max_{x_{p,i}, x_{s,i}} (x_{p,i} + x_{s,i}) \left(1 - \sum_{j=1}^2 (x_{p,j} + x_{s,j}) \right) - \frac{x_{p,i}^2 v}{\left(C_1 - \frac{\alpha}{2} \right)^2} - x_{s,i} \left(\frac{v \sum_{j=1}^2 x_{s,j}}{\alpha^2} \right). \quad (10)$$

Figures 9-12 depict how the equilibria in the variance-only model change as α increases for the case $C_1 = C_2 = 1$ and $v = 0.5$. Again, the monopolist outcomes are indicated by red asterisks. In all cases, both SPs will serve the same number of customers on their individual proprietary bands and on the shared band.

Figure 9 shows that delivered price initially rises with α and then begins to fall. This means that the volume of customers served initially falls and then rises. Since the firms are serving more people at lower prices, as α increases, the latency also increases causing a rise in prices. It is much easier to see the corresponding fall and rise of the total latency in Figure 12. Interestingly, the rise in latency as more sharing occurs is not as steep as seen in the mean-only Cournot model in Section IV-A. These changes indicate that as α increases, the firms are competing more. Figure 10 shows a slight rise and fall in profit and Figure 11 shows a rise in consumer surplus.

The initial drop in consumer surplus can be attributed to the fact that as α begins to increase, the price slightly increases and the latency decreases. However, as more customers are being served and experiencing a higher latency, the consumer surplus slightly increases. When consumers see lower latency, fewer customers are being served and prices will be higher. In this case we are able to observe the firms make use of the multiplexing gain.

To more closely examine the variation in the equilibrium profits and latency as α changes, we next examine the derivatives of these quantities versus α . Figure 13 graphs the derivative of equilibrium profit (denoted by π) with respect to α against α . Figure 13b graphs the derivative of equilibrium latency with respect to α . Note that both of these derivatives are near zero for all α .¹¹

In Figure 14, we show the impact of changing the variance on the equilibrium profit in the variance-only model. Figure 14a sets $\alpha = 0.5$. Figure 14b sets $\alpha = 0.75$. From these two graphs we are able to observe that the profit is relatively flat as variance increases, with a slight decreasing trend.

¹¹Analytically we can characterize these quantities for α close to zero, but omit this here due to space considerations.




Fig. 9: Delivered Price (\$) vs. α for variance-only model.




Fig. 12: Latency vs. α for variance-only model.





Fig. 10: Profit (\$) vs. α for variance-only model.



(a) $d\pi/d\alpha$ vs. α .





Fig. 11: Consumer Surplus (\$) vs. α for variance-only model.



(b) $dl/d\alpha$ vs. α .

Fig. 13: Derivatives of profit and latency with respect to α for the variance-only Cournot model.




Fig. 14: Firm profit versus variance v for the variance-only Cournot model with different choices of α .

V. CONCLUSIONS

We studied the impact of pooling spectrum with statistical multiplexing gains on the competition between SPs. We utilized a model based on Cournot competition with congestible resources, in which a mean-variance congestion cost was used to provide a tractable model that captured statistical multiplexing. Our results show that a monopolist would benefit from pooling spectrum through an increased profit. However, in a competitive market, the profit gains from pooling vanish due to the competition between the SPs. This suggests that SPs would not have an incentive to enter into a pooling agreement unless that agreement somehow also limited competition (e.g. by placing limits on the amount of traffic a SP could place on the shared band). We leave the study of such agreements to future work.

Pooling was shown to benefit consumer surplus as the increased competition resulted in lower prices, which suggests that from a policy point-of-view ways to encourage spectrum pooling might be desirable.

We did not address the impact of spectrum pooling on infrastructure costs. In some cases, e.g. if RAN sharing is utilized, pooling can reduce such costs, while in others costs could increase (e.g., due to the need for deploying some type of spectrum coordinator). Integrating such costs into this model is another direction for future work.

REFERENCES

- [1] T. A. Weiss and F. K. Jondral, "Spectrum pooling: An innovative strategy for the enhancement of spectrum efficiency," *IEEE Communications Magazine*, vol. 42, no. 3, pp. S8–14, 2004.
- [2] E. A. Jorswieck, L. Badia, T. Fahldieck, E. Karipidis, and J. Luo, "Spectrum sharing improves the network efficiency for cellular operators," *IEEE Communications Magazine*, pp. 129–136, 2014.
- [3] Y. Teng, Y. Wang, and K. Horneman, "Co-primary spectrum sharing for denser networks in local area," in *2014 9th International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CROWNCOM)*, 2014, pp. 120–124.
- [4] M. Rahman, S. H. Ahmed, and M. Yuksel, "Proof of sharing in inter-operator spectrum sharing markets," in *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2018, pp. 1–5.
- [5] R. Ratasuk, N. Mangalvedhe, and A. Ghosh, "LTE in unlicensed spectrum using licensed-assisted access," in *2014 IEEE Globecom Workshops (GC Wkshps)*, 2014, pp. 746–751.
- [6] Y.-D. Lin, H.-T. Chien, H.-W. Chang, C.-L. Lai, and K.-Y. Lin, "Transparent RAN sharing of 5G small cells and macrocells," *IEEE Wireless Communications*, vol. 24, no. 6, pp. 104–111, 2017.
- [7] F. P. Kelly, *Stochastic Networks: Theory and Applications* (Editors F.P. Kelly, S. Zachary and I.B. Ziedins), ser. Royal Statistical Society Lecture Notes Series. Oxford University Press, 1996, vol. 4, ch. Notes on effective bandwidths, pp. 141–168.
- [8] R. Berry, M. Honig, T. Nguyen, V. Subramanian, and R. Vohra, "The value of sharing intermittent spectrum," *Management Science*, vol. 66, no. 11, pp. 5242–5264, 2020.
- [9] P. Maillé and B. Tuffin, "Price war in heterogeneous wireless networks," *Computer Networks*, vol. 54, no. 13, pp. 2281–2292, 2010.
- [10] E. Kavurmacioglu, M. Alanyali, and D. Starobinski, "Competition in secondary spectrum markets: Price war or market sharing?" in *2012 IEEE International Symposium on Dynamic Spectrum Access Networks*, 2012, pp. 440–451.
- [11] T. Nguyen, H. Zhou, R. Berry, M. Honig, and R. Vohra, "The cost of free spectrum," *Operations Research*, vol. 64, no. 6, pp. 1217–1229, 2016.
- [12] C. Liu and R. A. Berry, "Competition with shared spectrum," in *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2014, pp. 498–509.
- [13] D. Acemoglu and A. Ozdaglar, "Competition and efficiency in congested markets," *Mathematics of Operations Research*, vol. 32, no. 1, pp. 1–31, February 2007.
- [14] F. Boccardi, H. Shokri-Ghadikolaei, G. Fodor, E. Erkip, C. Fischione, M. Kountouris, P. Popovski, and M. Zorzi, "Spectrum pooling in mmwave networks: Opportunities, challenges, and enablers," *IEEE Communications Magazine*, vol. 54, no. 11, pp. 33–39, 2016.
- [15] G. Middleton, K. Hooli, A. Tolli, and J. Lilleberg, "Inter-operator spectrum sharing in a broadband cellular network," *IEEE Ninth International Symposium on Spread Spectrum Techniques and Applications*, pp. 376–380, 2006.
- [16] W. Leher and N. Jesuale, "Spectrum pooling for next generation public safety radio systems," in *3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*, 2008, pp. 1–23.
- [17] R. Berry, M. Honig, T. Nguyen, V. Subramanian, H. Zhou, and R. Vohra, "On the nature of revenue-sharing contracts to incentivize spectrum-sharing," in *2013 Proceedings IEEE INFOCOM*, 2013, pp. 845–853.
- [18] B. Singh, S. Hailu, K. Koufos, A. A. Dowhuszko, O. Tirkkonen, R. Jäntti, and R. Berry, "Coordination protocol for inter-operator spectrum sharing in co-primary 5G small cell networks," *IEEE Communications Magazine*, vol. 53, no. 7, pp. 34–40, 2015.