The Economics of Shared Data Plans

Soumya Sen, Carlee Joe-Wong, Sangtae Ha
Princeton University, NJ, USA
{soumyas, cjoe, sangtaeh}@princeton.edu

Abstract

The demand for mobile broadband is doubling every year, forcing ISPs to use pricing as a congestion control mechanism. The largest US ISPs, AT&T and Verizon, have already terminated unlimited offerings in favor of "cap and meter" pricing plans with $10/GB overages. More recently, both ISPs have introduced shared data plans in which multiple devices share a data cap. However, such measures to curb demand can have adverse implications for mobile commerce and online content consumption. Hence, a study of the role of overage fees, usage caps, and shared data plans on the consumer's utility is needed. In this work, we introduce an analytical framework for studying the economics of such shared data plans and model the consumer utility of choosing between shared and separate (individual) data plans for their devices. We utilize usage data from a trial of 34 iPhone and iPad users to explore this tradeoff, and show that the choice between individual and shared data plans depends heavily on consumer's willingness to reduce usage upon exceeding the data cap. Our work creates a framework to study this dependency and indicates the importance of analyzing the impact of such data plans on consumer choice.

1. Introduction

The demand for mobile data is rising every year; the Cisco Visual Networking Index for 2012 projects an 18-fold increase in global mobile data traffic between 2011-2016. Mobile video is predicted to be the fastest-growing consumer mobile service, increasing from 271 million users in 2011 to 1.6 billion users in 2016 [Cisco VNI 2012]. Much of this demand growth is driven by the popularity of smart devices, bandwidth-hungry applications, cloud-based services, and media-rich web content. This adoption of smartphones and portable devices by consumers has also fuelled the popularity of mobile commerce [Kerschberg 2012], which is now projected to account for almost 25% of total e-commerce revenues by 2017 [ABI 2012]. In 2011, the mobile commerce market doubled in size to $65.6 Billion, partly due to increases in adoption and consumption by smartphone users in both mature and developing markets [ABI 2012].

However, the sustainability of this growth in the consumption of online content and mobile commerce is increasingly under threat from Internet Service Providers' (ISP) inability to manage the capacity of their wireless networks to accommodate the increasing traffic volume. The major US ISPs, like AT&T, Verizon, and Comcast, have abandoned their traditional unlimited "flat rate" data plans in favor of throttling, tiered usage-based (metered) pricing, and data caps [Sen 2012]. More recently, AT&T and Verizon introduced shared data plans (AT&T's Mobile Share and Verizon's Share Everything) for families and users with multiple devices to share their limited quota across devices. Verizon has also replaced individual plans by shared data plans [Chen 2012]. Hence, understanding how shared data plans can help or hurt consumers is crucial in sustaining demand for mobile data consumption and m-commerce activities. We therefore consider the issue from a consumer's perspective, although similar economic modeling from an ISP's perspective is also equally important, and will be addressed in future works.

In this paper, we introduce a simple analytical framework to model the tradeoffs between choosing a shared data plan and individual data plans for different devices from a consumer perspective. In particular, we show that even in this simplified setting, this choice is non-trivial for a consumer, with a complex dependency on users' psychology of self-censoring their usage upon exceeding their data caps. Restricting the consumer's choice by eliminating individual device data plans can therefore hurt certain types of consumers. Although many shared data plans currently provide free texting and talk time, VoIP has effectively made these services free. Hence, we focus only on data usage caps, which directly impact the Internet ecosystem.

2. Model

We develop a framework to understand the consumer utility and the trade-offs in choosing between individual data plans and shared data plans for his/her mobile devices. Specifically, our model considers two devices with different marginal values for bandwidth consumed, v_1 and v_2, and different amounts of maximum usage, C_1^{max} and C_2^{max}, which can be either put on a shared plan or individual plans by their owner. The nature of the devices under consideration can determine the relative values of these parameters, as we discuss next.

In the first example scenario, consider two devices in a family with two different owners who have different priorities (*e.g.,* a parent and a child). Typically, the value for each unit of bandwidth that a smart device consumes is lower for a child than a parent ($i.e., v_1 < v_2$), and so is the data cap on a child's device ($C_1^{max} < C_2^{max}$). In the second setting, consider a single user with two devices (*e.g.,* iPhone and iPad): one is a personal productivity device for wireless data access (*e.g.,* iPhone) and the other is for entertainment (*e.g.,* iPad). The value of each unit of bandwidth consumed on the productivity device, v_1, is higher than that of the entertainment device, v_2 ($i.e., v_1 > v_2$), but the latter has a higher maximum data demand ($C_1^{max} < C_2^{max}$)[1].

In both these scenarios, a monopolist ISP is assumed to offer a continuum of usage-based pricing plans with different bandwidth caps from which consumers can choose[2]. If a consumer exceeds the chosen bandwidth cap, overage fees are charged. Mathematically, an ISP's individual base plans are priced at p_1 and p_2 per unit of bandwidth consumed (measured, *e.g.*, in MB) for the two types of devices, respectively, and a shared data plan of p_s per unit of bandwidth consumed on either device ($p_s > p_1, p_2$). Additionally, the shared plan also has a fixed monthly fee g_1 and g_2 for each of the two devices (*e.g.,* iPads and iPhones have different flat fees for being on the shared plan). For instance, AT&T charges $85 for 1GB of data per month for one smartphone. Adding a smartphone to this shared plan costs $45, with an overage charge of $15/GB [Molen 2012]. For the case of individual data plans for the two devices, a consumer chooses monthly bandwidth caps of B_1 and B_2 and can incur overage fees of o_1 and o_2 for each unit of bandwidth consumed above these caps. Under the shared data plan, a consumer chooses a bandwidth cap of B_s across the two devices and incurs an overage cost of o_s per unit of extra bandwidth consumed on either device ($o_s > o_1, o_2$).

Because of the overage fees involved, we assume that for individual device data plans, when the device with a lower value (*e.g.,* a child's smartphone or the iPad, in the two examples above) exceeds the consumer's chosen cap for that device, then the user continues to consume data but cuts back on his demand and consumes only a fraction $(1 - \alpha)$ of the demand over the cap. This typically happens because overage warnings and higher fees trigger the user psychology of self-restraint. Similarly, for the shared case, when the cap is exceeded, the user curtails the usage on the lower-value device by a factor α. Due to space constraints we only perform the analysis for the case $v_1 < v_2$ and $C_1^{max} < C_2^{max}$ (*i.e.,* a child's and parent's devices); the analysis for other scenarios is similar.

---

[1]Typically smartphones are actively used by consumers for their work-related data access and consume about 450 MB/month on an average, whereas users consume roughly 3.6 GB on iPads by streaming music and mobile videos.

[2]We assume usage-based pricing for simplicity, but tiered data plans can be considered using a similar framework. AT&T's tiered plans of $30 for 3GB, $50 for 5GB, already approximates *usage-based* plans of $10/GB. We consider a continuum of data cap choices for users (like a continuum of tiered plans over a range of data cap options), but the model can be easily discretized.

The consumer's objective is to choose between shared or individual plans and an optimal data cap. This decision involves three stages: first a consumer needs to choose the type of plan (shared or individual) and the data cap, then the demand on the two devices is realized, and third, if the demand exceeds the cap, then the consumer react by reducing some usage from the device with lesser value due to the higher overage fees. This sequential decision process is shown in Fig. 1. The problem needs to be solved backward by computing the expected utility for a given choice of data cap, and then choosing the cap and plan that maximizes the consumer's utility.

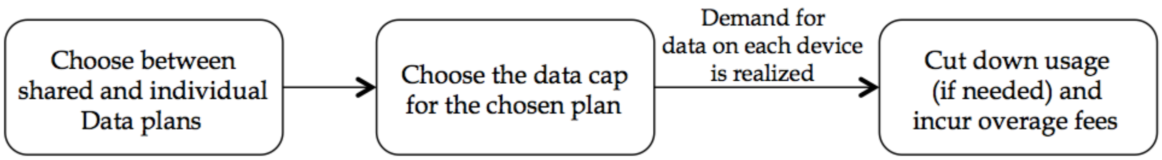

Figure 1: The stages of sequential decision process for a data plan consumer.

3. Analysis

We suppose that the consumer's realized demand for device i is c_i, which follows a distribution $f_i(c_i)$ that varies from *0* to C_i^{max}. The utility derived by a consumer from consuming some amount c_i is assumed to be a concave function of data consumed, $v_i \log(1 + c_i)$. For the individual data plan of device 1, the consumer's utility for choosing a data plan of cap B_1 is:

$$U_1 = \begin{Bmatrix} v_1 \log(1 + c_1) - p_1 B_1, & c_1 < B_1 \\ v_1 \log(1 + c_1 - \alpha(c_1 - B_1)) - p_1 B_1 - o_1(1 - \alpha)(c_1 - B_1), & c_1 \geq B_1 \end{Bmatrix}$$

Note that upon exceeding the cap on device 1, a fraction of excess demand α is lost.

For device 2's individual plan, the utility for the consumer for choosing a data plan of cap B_2 is:

$$U_2 = \begin{Bmatrix} v_2 \log(1 + c_2) - p_2 B_2, & c_2 < B_2 \\ v_2 \log(1 + c_2) - p_2 B_2 - o_2(c_2 - B_2), & c_2 \geq B_2 \end{Bmatrix}$$

Under individual plans, a consumer's utility from choosing a data plan with cap B_1 for device 1 is

$$E(U_1|B_1) = \int_0^{B_1} [v_1 \log(1 + c_1) - p_1 B_1] f_1(c_1) dc_1 + \int_{B_1}^{C_1^{max}} [v_1 \log(1 + c_1 - \alpha(c_1 - B_1)) - p_1 B_1 - o_1(1 - \alpha)(c_1 - B_1)] f_1(c_1) dc_1 \quad (1)$$

Setting the derivative $dE(U_1|B_1)/dB_1 = 0$, the optimal data plan cap for consumer's device 1, B_1^*, and the resulting utility, $U_1^* = E(U_1|B_1^*)$, can be estimated. A similar calculation for $U_2^* = E(U_2|B_2^*)$ gives the individual data plan with the cap that is optimal for device 2.

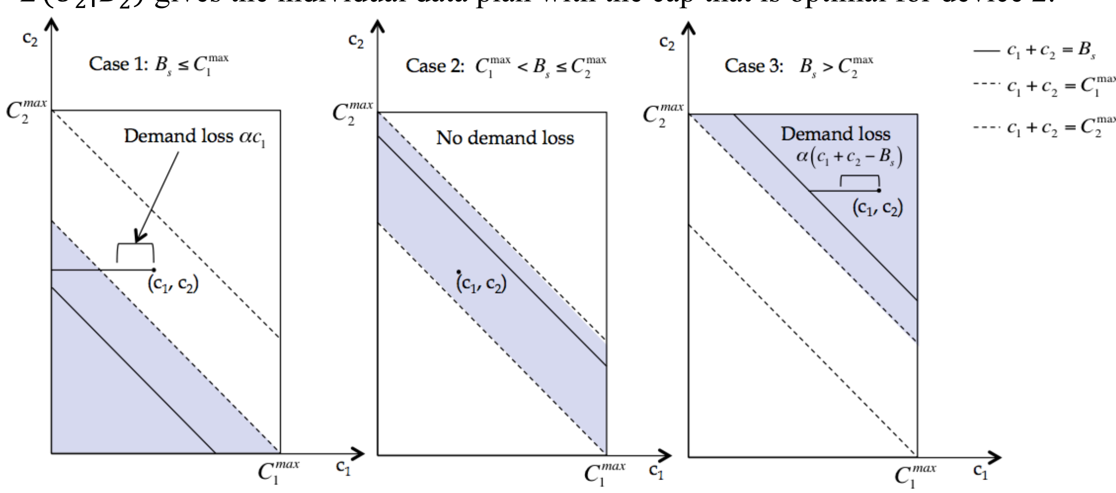

Figure 2: The three regions for which the optimal B_s need to be calculated. The one with the highest $E(U|B_s)$ is chosen as the optimal data cap for a consumer.

For the shared data plan, if the demand across the devices exceeds the shared cap B_s, then the lower priority device 1 curbs its realized usage c_1 by a fraction α of the amount by which it exceeds the shared cap B_s. As in the case of individual data plans, the utility of a user for a chosen value of data cap B_s needs to be computed by integrating over the demand distribution. But the limits of the integration and the utility derived are different depending on the value of the chosen B_s compared to C_1^{max} and C_2^{max}. As shown in Figure 2, there are 3 such regions in the (C_1^{max}, C_2^{max})-plane for which the expected utility needs to be calculated and the value of B_s^* that maximizes it should be chosen as the data cap for the user's shared plan.

If $c_1 + c_2 \leq B_s$, the shared data cap is not exceeded and the user experiences no loss in demand. If $c_2 \leq B_s$ but $c_1 + c_2 \geq B_s$, device 1 experiences a loss in demand proportional to the amount by which the cap B_s is exceeded, i.e., a loss of $\alpha(c_1 + c_2 - B_s)$; if, however, $c_2 > B_s$ and $c_1 + c_2 \geq B_s$, device 1 experiences a demand loss proportional only to the amount by which usage of device 1 exceeds the cap, i.e., αc_1. The expected utility for $B_s < C_1^{max}$ is thus

$$\begin{aligned}
E\left(U_s \mid B_s < C_1^{\max}\right) = &\int_0^{B_s} \int_0^{B_s - c_2} \left[v_1 \log\left(1 + c_1\right) + v_2 \log\left(1 + c_2\right) - p_s B_s\right] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2 \\
&+ \int_0^{B_s} \int_{B_s - c_2}^{C_1^{\max}} \Big[v_1 \log\left(1 + (1-\alpha)c_1 - \alpha\left(c_2 - B_s\right)\right) + v_2 \log\left(1 + c_2\right) - p_s B_s \\
&\qquad - o_s(1-\alpha)\left(c_1 + c_2 - B_s\right)\Big] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2 \\
&+ \int_{B_s}^{C_2^{\max}} \int_0^{C_1^{\max}} \Big[v_1 \log\left(1 + (1-\alpha)c_1\right) + v_2 \log\left(1 + c_2\right) - p_s B_s \\
&\qquad - o_s\left((1-\alpha)c_1 + c_2 - B_s\right)\Big] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2
\end{aligned}$$

Similarly, we compute

$$\begin{aligned}
E\left(U_s \mid C_1^{\max} < B_s < C_2^{\max}\right) = &\int_0^{B_s} \int_0^{B_s - c_2} \left[v_1 \log\left(1 + c_1\right) + v_2 \log(1 + c_2) - p_s B_s\right] f_1(c_1) f_2(c_2)\ dc_1\ dc_2 \\
&+ \int_0^{B_s - C_1^{\max}} \int_0^{C_1^{\max}} \left[v_1 \log(1 + c_1) + v_2 \log\left(1 + c_2\right) - p_s B_s\right] f_1(c_1) f_2(c_2)\ dc_1\ dc_2 \\
&+ \int_{B_s - C_1^{\max}}^{C_1^{\max}} \int_{B_s - c_2}^{C_1^{\max}} \Big[v_1 \log\left(1 + (1-\alpha)c_1 - \alpha\left(c_2 - B_s\right)\right) + v_2 \log\left(1 + c_2\right) - p_s B_s \\
&\qquad - o_s(1-\alpha)\left(c_1 + c_2 - B_s\right)\Big] f_1(c_1) f_2(c_2)\ dc_1\ dc_2 \\
&+ \int_{B_s}^{C_2^{\max}} \int_0^{C_1^{\max}} \Big[v_1 \log\left(1 + (1-\alpha)c_1\right) + v_2 \log\left(1 + c_2\right) - p_s B_s \\
&\qquad - o_s\left((1-\alpha)c_1 + c_2 - B_s\right)\Big] f_1(c_1) f_2(c_2)\ dc_1\ dc_2
\end{aligned}$$

$$\begin{aligned}
E\left(U_s \mid C_2^{\max} < B_s < C_1^{\max} + C_2^{\max}\right) = &\int_0^{B_s - C_1^{\max}} \int_0^{C_1^{\max}} \left[v_1 \log\left(1 + c_1\right) + v_2 \log\left(1 + c_2\right) - p_s B_s\right] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2 \\
&+ \int_{B_s - C_1^{\max}}^{C_2^{\max}} \int_0^{B_s - c_2} \left[v_1 \log\left(1 + c_1\right) + v_2 \log\left(1 + c_2\right) - p_s B_s\right] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2 \\
&+ \int_{B_s - C_1^{\max}}^{C_2^{\max}} \int_{B_s - c_2}^{C_1^{\max}} \Big[v_1 \log\left(1 + (1-\alpha)c_1 - \alpha\left(c_2 - B_s\right)\right) + v_2 \log\left(1 + c_2\right) - p_s B_s \\
&\qquad - o_s(1-\alpha)\left(c_1 + c_2 - B_s\right)\Big] f_1\left(c_1\right) f_2\left(c_2\right)\ dc_1\ dc_2.
\end{aligned}$$

These equations may then be individually solved for the optimal budgets B_s^{*i}, $i = 1,2,3$ corresponding to the three cases above. The expected utility values at these optimal budgets can then be compared to find $E(U_s|B_s^*)$ for the shared data plan. Similar methods can be applied if $v_1 > v_2$ (e.g., iPhones and iPads).

4. Data and Simulations

We now apply the analytical framework introduced above to empirical data. We first estimate the distributions of the c_1 and c_2 variables, i.e., the distribution of the monthly usage volume for devices 1 and 2. As shared mobile data plans have come onto the market only recently, we use data from devices on individual data plans. Our data comes from 19 iPhone (i.e., device 1) and 15 iPad (i.e., device 2) users. We recruited trial participants from fourteen different academic and administrative divisions of our university, as well as their family members, and measured their 3G usage at an hourly granularity for three months. Based on this usage data, our trial participants fell into two groups: high usage (over 900MB for iPad users, and over 400MB for iPhone users) and low usage. We used maximum likelihood estimation to fit a β distribution to the monthly usage data for users of each device, in each group.[3] We obtain similar β distribution parameter values for the different groups of users: $(a, b) = (0.482, 0.549)$ for low-usage iPhone users, (0.344, 0.224) for high-usage iPhone users, (0.384, 0.408) for low-usage iPad users, and (0.515, 0.191) for high-usage iPad users. In the simulations below, we utilize the parameter values for the low usage groups of iPhone (device 1) and iPad (device 2) users; similar results may be obtained with the parameters for high-usage users.

We suppose that the user values iPhone over iPad usage and that iPad demand is lost over the data cap. Figure 3 shows the optimal budgets for the shared and individual data plans with their corresponding expected utility values, for a given set of marginal valuations and prices. If the user is not willing to cut back on her over-the-cap usage ($\alpha < 0.1$), then she prefers an individual data plan: since the shared plan's marginal price $p_s = 14$ is more expensive than the individual plans' ($p_1 = 8, p_2 = 10$), her optimal cap on a shared data plan does not account for large usage amounts on both devices. Similarly, if the user reduces her iPad's usage over the cap ($\alpha > 0.6$), she prefers individual data plans; most over-the-cap usage will take place on device 1, which can be accounted for with device 1's data cap. The expected utility of both individual and shared data plans is thus non-monotonic, reflecting the tradeoff between the utility of consuming more and disutility of exceeding the data cap.

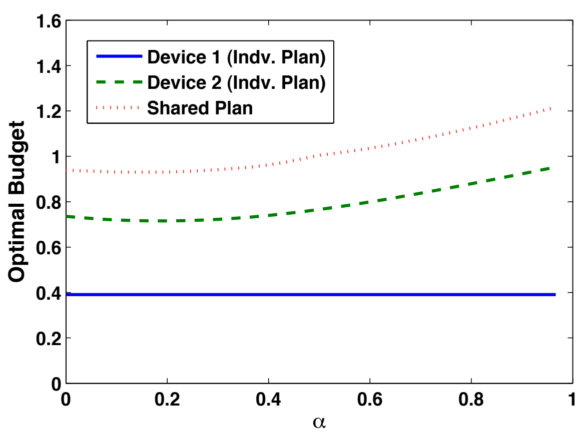

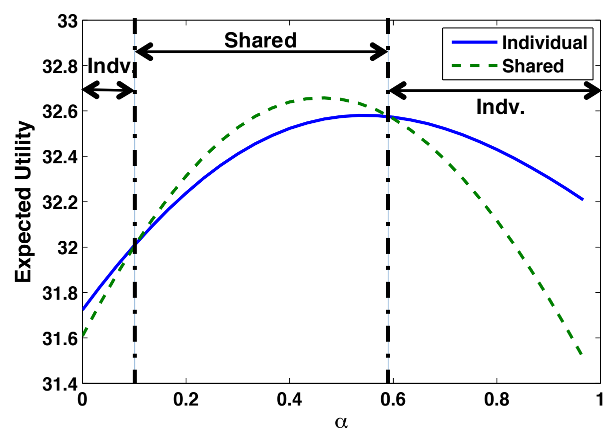

Figure 3: Optimal budgets and expected utility with β distribution parameters from real data. Prices and valuations (\$/GB) are $v_1 = 60, v_2 = 40, p_1 = 8, p_2 = 10, o_1 = 15, o_2 = 18, p_s = 14, o_s = 20$; maximum demand is $C_1^{max} = 1$ GB, $C_2^{max} = 2$ GB. Fixed fees for the shared data plans are (in \$) $g_1 + g_2 = 2$.

[3] Though our model assumes a demand distribution for one user of each device, we aggregate data from a group of users in order to obtain a sufficient number of monthly usage data points to fit a distribution.

Qualitatively similar behavior, with the user switching between preferring shared and individual data plans, can be obtained for other demand distributions (e.g., a uniform distribution) and for the case in which demand is lost from device 1.

Related Work

The question of choosing between shared and individual data plans has some parallels to the decision in manufacturing systems in provisioning between flexible and dedicated resources in presence of demand uncertainty [Fine 1990, Van Mieghem 1998]. There are two key differences between this paper and these models. First, rather than exploring the benefits of hedging against uncertainty with shared resources; the focus of this work is to understand the impact of the fee structure on the data plan choices users make. Secondly, in manufacturing systems, excess demand is lost due to time lag of ramping up capacity; but in case of shared data plans, users can reduce their usage in response to exceeding the data caps. As we show in this work, the level of willingness of a user to curb down on their excess usage can not only affect the data caps they choose a priori but also the very choice (shared or individual) of data plans for their devices.

Conclusions

In this work, we introduce an analytical framework to study user choice between individual and shared data plans for their mobile devices. In particular, we study the impact of overage fees on users' decisions of whether to adopt individual or shared data plans. We use empirical usage data from iPhone and iPad users to demonstrate that user choice of shared or individual data plans depends non-trivially on their psychological response in reducing usage upon exceeding their monthly data cap. Users who are either very willing or very unwilling to reduce usage prefer individual plans, while those between these extremes prefer shared plans. Thus, the economic impact of shared data plans deserves future study: their effect on users' mobile data consumption and the Internet ecosystem, can be non-intuitive. Our work provides an initial framework and results in this direction. Future directions for investigation include considering the issue from an ISP's perspective of how to choose between plans with different usage and cap structures, how to price them, and the social welfare realized under the profit maximizing decisions of the ISP.

References

[Cisco VNI 2012] Cisco Visual Networking Index White Paper. February, 2012. Available online at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf

[Kerschberg 2012] "How Instant Mobile Commerce Will Disrupt The Way You Do Business," B. Kerschberg, Forbes Tech, Online, July 2012.

[ABI 2012] "M-Commerce Growing from 24% of Total E-Commerce Market Value on Back of Smartphone Adoption, Claims ABI Research," ABI Research Press Release, Online, July 30, 2012.

[Chen 2012] "Shared Mobile Data Plans: Who Benefits?" B. X. Chen, The New York Times, July 19, 2012.

[Fine 1990] "Optimal investment in product-flexible manufacturing capacity," C. H. Fine and R. M. Freund, Management Science, 36(4), 1990, pp. 449-465.

[Molen 2012] "AT&T and Verizon shared data plans compared," B. Molen, Online, EnGadget, July 2012.

[Sen 2012] "Incentivizing Time-Shifting of Data: A Survey of Time-Dependent Pricing for Internet Access", S. Sen, C. Joe-Wong, S. Ha and M. Chiang, IEEE Communications Magazine, November, 2012.

[Van Mieghem 1998] "Investment strategies for flexible resources," J. A. Van Mieghem, Management Science, 44(8), 1998, pp. 1071-1078.