Generalizations of the Mitra-Wang Model of ISP Behavior in the Presence of Managed Service

By: Austin Ebel
ELEN 6767 - Prof. Debasis Mitra

1 Overview

A pressing question in the debate for Net Neutrality is how ISPs will react in the presence of *both* Best Effort and Managed Service pricing schemes. The current Best Effort service can be considered *free* at the expense of decreased quality of service (QoS) in the form of delay. Managed Service, on the other hand, provides guaranteed QoS, with no delay, at some per-use fee set by the ISP. Mitra and Wang [1] have previously modeled the interplay between these two services and highlight that, while Best Effort does not generate direct revenue, it has long been recognized as the driving force for the invention of new applications that creates new demand for Managed Service.

In this paper, they assume a uniform distribution of maximum willingness to pay (WtP), or willingness to tolerate delay, for all users N who share a network link. In practice, these distributions can vary, and will likely follow a Gaussian-like distribution - in which the rate of subscribing to a product begins to decrease past some set price [2-4]. Here, I present an extension of Mitra and Wang's model to arbitrary Beta distributions, Beta(α, β), which can be shown to encapsulate their original $U\{0,1\}$ WtP assumption. Furthermore, I will look at how varying distribution parameters α, β, in conjunction with varying per-use *Managed Service* fees, affect steady-state ISP *profits*.

2 Prior Work

This section will attempt to summarize mathematical concepts found in the Mitra-Wang model. It will start by defining parameters found in the model, followed by key derivations that lead to a profit function used in this paper.

- **A**: the number of network applications.
- **N**: the number of customers who share a common resource (network link, etc.).
- **p**: the per-use *Managed Service* fee.
- **D**: the mean delay seen by *Best-Effort* users.
- C^{BE}: the bandwidth provisioned for Best Effort users.
- θ: user's index of willingness to pay for Managed Service (inability to tolerate delay).
- ω: the maximum amount of utility that a customer can derive by using the application an infinite number of times ($\lambda \to \infty$).
- γ: the rate of decay of marginal utility.
- δ: the birth rate of new applications.
- μ: the death rate of existing applications.
- κ: the maximum cost per unit of delay.
- η^{BE}: the slope of the Best Effort cost function.
- η^{MS}: the slope of the Managed Service cost function.
- α: scaling factor used in Managed Service cost function.

1. The utility gained from A applications used at a rate λ is,

$$u(\lambda) = A\omega(1 - e^{-\gamma\lambda}) \tag{1}$$

2. Customers choose the optimal usage level to maximize their net surplus,

$$\max_{\lambda}\{u(\lambda) - xA\lambda\} \tag{2}$$

   where $x = p$ in the *Managed Service* case, and $x = \theta\kappa D$ in the *Best Effort* case, resulting in an optimal usage rate of,

$$\lambda^* = \frac{1}{\gamma} ln \frac{\omega\gamma}{x} \tag{3}$$

3. Let θ_b be the critical willingness to pay for the user who is indifferent to either services. Thus,

$$\theta_b = \frac{p}{\kappa D} \tag{4}$$

4. The total *Managed Service* usage rate is,

$$\begin{aligned} \Lambda^{MS} &= AN\left(\int_{\theta_b}^{1} d\theta\right)\lambda^* \\ &= AN(1-\theta_b)\lambda^* \\ &= AN\tfrac{1-\theta_b}{\gamma} ln \tfrac{\omega\gamma}{p} \end{aligned} \tag{5}$$

5. The total *Best Effort* usage rate is,

$$\begin{aligned} \Lambda^{BE} &= \tfrac{AN}{\gamma}\int_0^{\theta_b} ln \tfrac{\omega\gamma}{\theta\kappa D} d\theta \\ &= \tfrac{AN\theta_b}{\gamma}(1 + ln \tfrac{\omega\gamma}{\theta_b \kappa D}) \\ &= \tfrac{AN\theta_b}{\gamma}(1 + ln \tfrac{\omega\gamma}{p}) \end{aligned} \tag{6}$$

6. Mean delay is given by the M/M/1 formula,

$$D = \frac{1}{C^{BE} - \Lambda^{BE}} \tag{7}$$

7. Profit rate is defined as,

$$\Pi(p, C^{BE}) = p\Lambda^{MS} - \eta^{BE}C^{BE} - \eta^{MS}(A)^{\alpha}N(1-\theta_b)\frac{1}{\gamma} ln \frac{\omega\gamma}{p} \tag{8}$$

8. Using Eq. (4), (6), (7), $\theta_b \leq 1$ and solving for θ_b,

$$\theta_b = \min\left[1, C^{BE}\left(\frac{AN}{\gamma}\left(1 + ln\frac{\omega\gamma}{p}\right) + \frac{\kappa}{p}\right)^{-1}\right] \tag{9}$$

9. Assume dynamic behavior of the model driven by the *Best Effort* service following,

$$\frac{dA}{dt} = \delta\Lambda^{BE} - \mu A \tag{10}$$

   In this case, a θ_b (and thus a D and C^{BE}) is forced in long-term equilibrium,

$$\theta_b = \frac{\mu\gamma}{N\delta}(1 + ln\frac{\omega\gamma}{p})^{-1} \tag{11}$$

$$D = \frac{p}{k\theta_b} = \frac{p}{k}\frac{N\delta}{\mu\gamma}(1 + ln\frac{\omega\gamma}{p}) \tag{12}$$

$$C^{BE} = \Lambda^{BE} + \frac{1}{D} = \frac{\mu}{\gamma}\left(A + \frac{\kappa\gamma}{Np}\left(1 + ln\frac{\omega\gamma}{p}\right)^{-1}\right) \tag{13}$$

3 Assertion and Validation

Mitra and Wang assume a uniform distribution of willingness to pay for the Managed Service of the form below:

$$f(w) = \begin{cases} 0 & w < 0,\ w > 1 \\ 1 & 0 \leq w \leq 1 \end{cases}$$

This work extends their uniform willingness to pay assumption to the following equation in order to understand how non-uniform willingness to pay affects total ISP profits,

$$f(w,\alpha,\beta) = \begin{cases} 0 & w < 0,\ w > 1 \\ \dfrac{1}{B(\alpha,\beta)}w^{\alpha-1}(1-w)^{\beta-1} & 0 \leq w \leq 1 \end{cases}$$

where,

$$B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)} \tag{14}$$

and,

$$\Gamma(n) = (n-1)! \tag{15}$$

Specific PDFs and CDFs corresponding to varying α and β values can be seen in Figure 1. This, in turn, directly changes the Λ^{BE} and Λ^{MS} required for long-term equilibrium, which percolates down to change most solved quantities in Eq. (5)-(13). Let us first look at how to modify Λ^{BE}.

What was previously a uniform distribution modeled by the implicit 1 in $\int_{\theta_b}^{1} d\theta$ must now be modified as such,

$$\Lambda^{BE} = AN\left(\int_0^{\theta_b}\left(ln\tfrac{\omega\gamma}{\theta\kappa D}\right)\theta^{\alpha-1}(1-\theta)^{\beta-1}d\theta\right)\lambda^* \tag{16}$$

which does not produce closed form solutions in the general case. This is a variant of the incomplete beta distribution [5], where the integral must be written as an infinite sum of the following form:

$$Beta(w_{\hat{i}},\alpha,\beta) = w_{\hat{i}}^{\alpha}\sum_{n=0}^{\infty}\frac{(1-b)_n}{n!(a+n)}w_{\hat{i}}^{n} \tag{17}$$

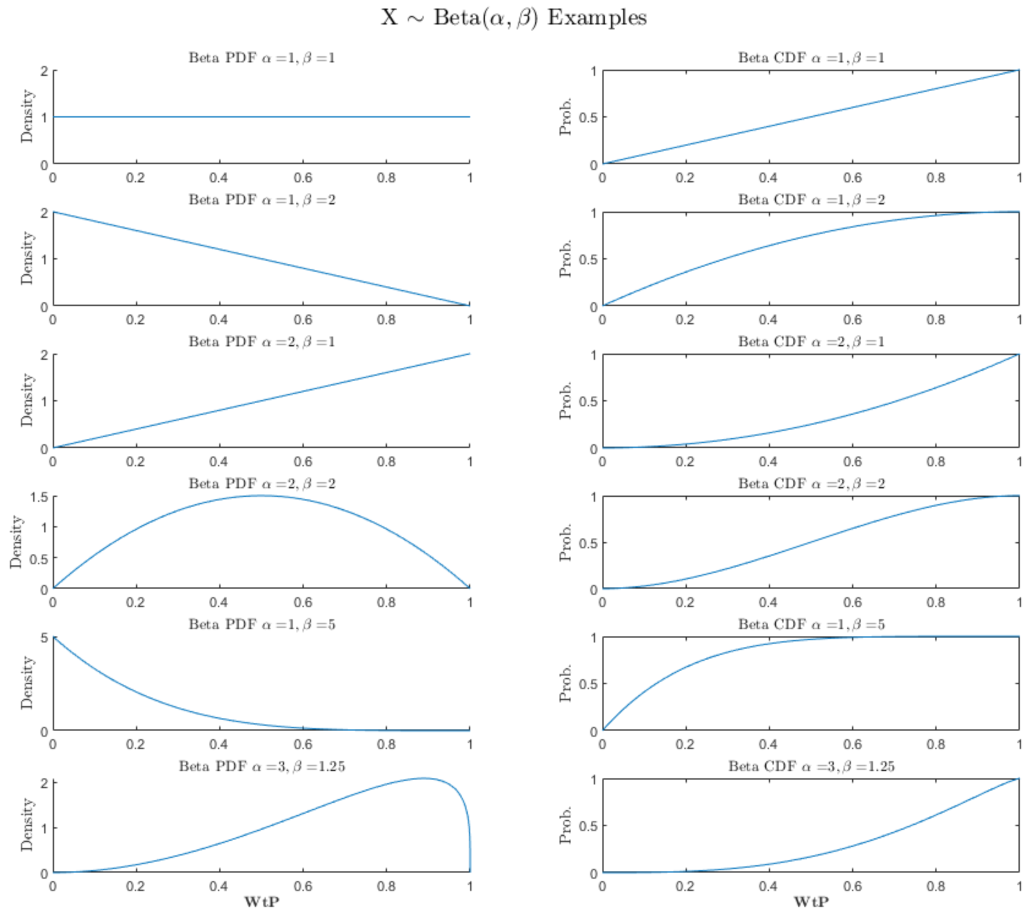

Figure 1: Beta Distribution PDF/CDF Examples

where $(1-b)_n$ is a Pochhammer symbol [6]:

$$(1-b)_n = \frac{\Gamma(1-b+n)}{\Gamma(1-b)} \tag{18}$$

So, while there is no closed form solution in the general case, this work will look at individual values of α, β that produce closed form solutions. Similarly,

$$\Lambda^{MS} = AN\left(\int_{\theta_b}^{1} \theta^{\alpha-1}(1-\theta)^{\beta-1} d\theta\right)\lambda^{*} \tag{19}$$

which again forms a variant of the incomplete beta distribution. That said, for specific values of α, β, solutions can be readily found. As is seen in the equation for the Beta distribution, setting $\alpha = 1$ and $\beta = 1$ will result in the uniform distribution. As an example, let us set $(\alpha, \beta) = (2, 2)$, and solve for Λ^{MS} and Λ^{BE}.

$$\Lambda^{MS} = AN\left(\int_{\theta_b}^{1} \theta(1-\theta)d\theta\right)\lambda^* \\ = \frac{AN}{6\gamma}(2\theta_b^3 - 3\theta_b^2 + 1)ln\frac{\omega\gamma}{p} \tag{20}$$

$$\Lambda^{BE} = AN\left(\int_{0}^{\theta_b}(ln\frac{\omega\gamma}{\theta\kappa D})\theta(1-\theta)d\theta\right)\lambda^* \\ = \frac{AN}{\gamma}\left[1.5\theta_b^2\left(2ln\frac{\omega\gamma}{p}+1\right) - 0.667\theta_b^3\left(3ln\frac{\omega\gamma}{p}+1\right)\right] \tag{21}$$

Given,

$$\Lambda^{BE} = \frac{\mu}{\gamma}A \tag{22}$$

a similar derivation, albeit incredibly complicated, can be done to solve for θ_b (and thus D and C^{BE}) in this new general case.

For the rest of this section, let us look at how varying α and β (and thus varying the distribution of this willingness to pay) affects ISP profits. Recall,

$$\Pi(p, C^{BE}) = p\Lambda^{MS} - \eta^{BE}C^{BE} - \eta^{MS}(A)^{\alpha}N(1-\theta_b)\frac{1}{\gamma}ln\frac{\omega\gamma}{p} \tag{23}$$

Mitra and Wang assign values to the model parameters in a specific case as follows,

- $\omega = 0.5$
- $N = 500$
- $\frac{\mu}{\delta} = 50$
- $\eta^{BE} = 4000$
- $\eta^{MS} = 5000$
- $\alpha = 1.4$
- $\kappa = 0.5$

In order to replicate their previous work in profit maximization, I've chosen,

- $A = 1200$
- $\gamma = 2$
- $0.8 \leq p \leq 2$

Results from this can be seen in Figure 2. Thus, it appears that strictly decreasing PDFs (giving the effect of *increasing* the appeal of Best Effort service) will lead to more profits. This ultimately makes sense as one of the driving forces behind demand for Managed Service is demand for Best Effort.

4 Conclusion

This work set out to understand how willingness to pay distributions (willingness to tolerate the delay of Best Effort) outside of the uniform distribution affect steady-state ISP profits. By using the Beta distribution, this work explores many such PDFs with varying degrees of skew by modifying the α and β parameters. From this, we gain more key insights into the interplay between Best Effort and Managed Service - a key issue in the debate for Net Neutrality.

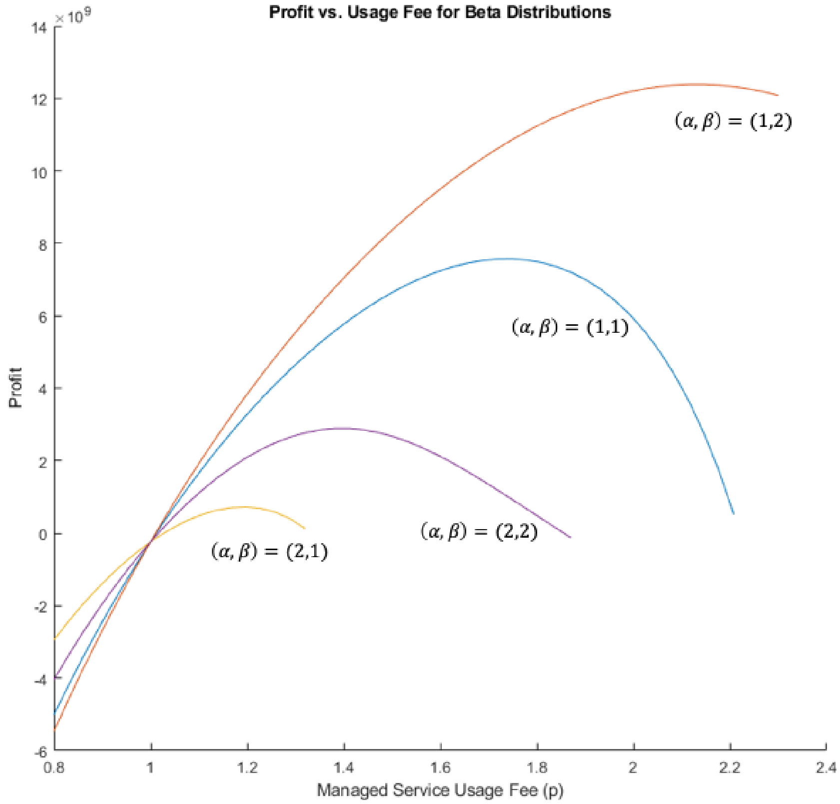

Figure 2: Profit vs. Usage Fee for Beta Distributions

References

[1] Mitra, Debasis & Wang, Qiong. (2015). Preservation of Best-Effort Service on the Internet in the Presence of Managed Services and Usage-Generated Applications.

[2] Das, Chhandita, et al. “Estimating Distributions of Willingness to Pay for Heterogeneous Populations.” Southern Economic Journal, vol. 75, no. 3, 2009, pp. 593–610. JSTOR, www.jstor.org/stable/27751406.

[3] Richard T. Carson, Mikołaj Czajkowsk, "A new baseline model for estimating willingness to pay from discrete choice models" Journal of Environmental Economics and Management.

[4] Riccardo Scarpa, Mara Thiene, Kenneth Train, Utility in Willingness to Pay Space: A Tool to Address Confounding Random Scale Effects in Destination Choice to the Alps, American Journal of Agricultural Economics, Volume 90, Issue 4, November 2008, Pages 994–1010, https://doi.org/10.1111/j.1467-8276.2008.01155.x

[5] https://mathworld.wolfram.com/IncompleteBetaFunction.html

[6] https://mathworld.wolfram.com/PochhammerSymbol.html