Competitive Provisioning of Online Services

Jayakrishnan Nair
Electrical Engineering, IIT Bombay

Joint work with Vijay Subramanian, Adam Wierman & Bert Zwart

facebook

twitter

digg

flickr

YouTube

Blogger

YAHOO!

Google

WIKIPEDIA
The Free Encyclopedia

Characteristics of online services

1. Majority are offered for free (and supported by Ads.)
   - Internet advertising revenue totaled $26 billion in 2010
2. Users are highly congestion (delay) sensitive
   - Google search: 0.5 seconds additional delay => 20% drop in traffic
3. <u>Positive network effects</u> in the user base
   - **Users derive utility from other people using the service**
4. High level of competition between providers

Focus: How these factors lead to capacity provisioning by profit maximizing firms

Case of a single provider

Model

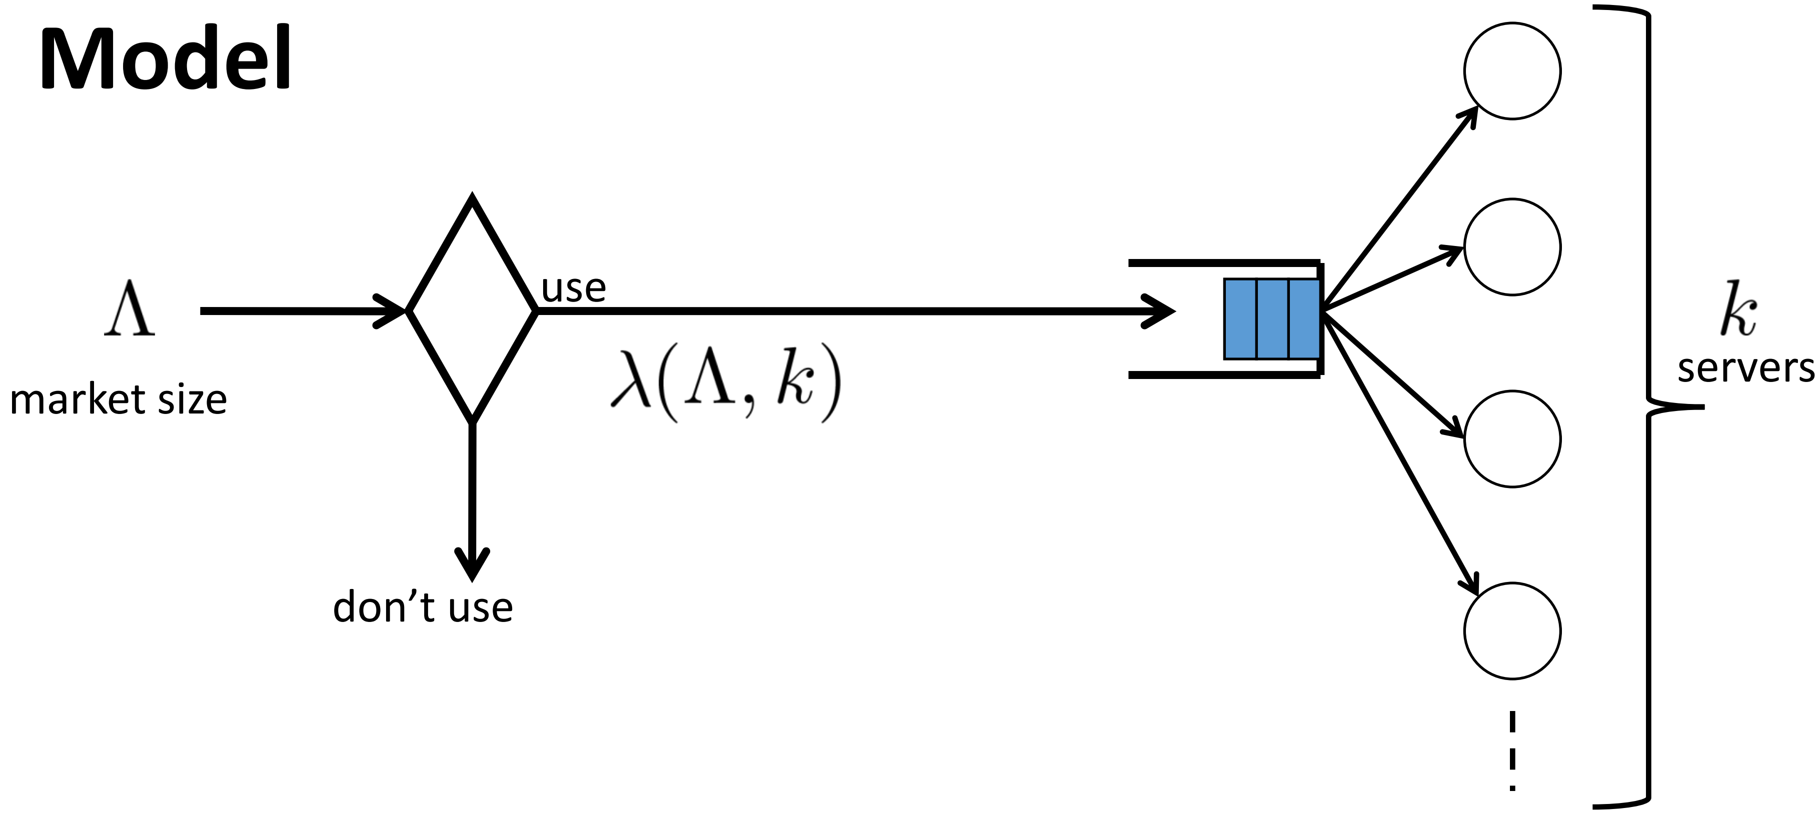

For mathematical tractability, we assume:

- Poisson arrivals
- Exponential service times
- FCFS service

} $\mathrm{M/M/}k$

WLOG, assume mean service time = 1 $\Rightarrow \lambda(\Lambda, k) < k$

Model (user base)

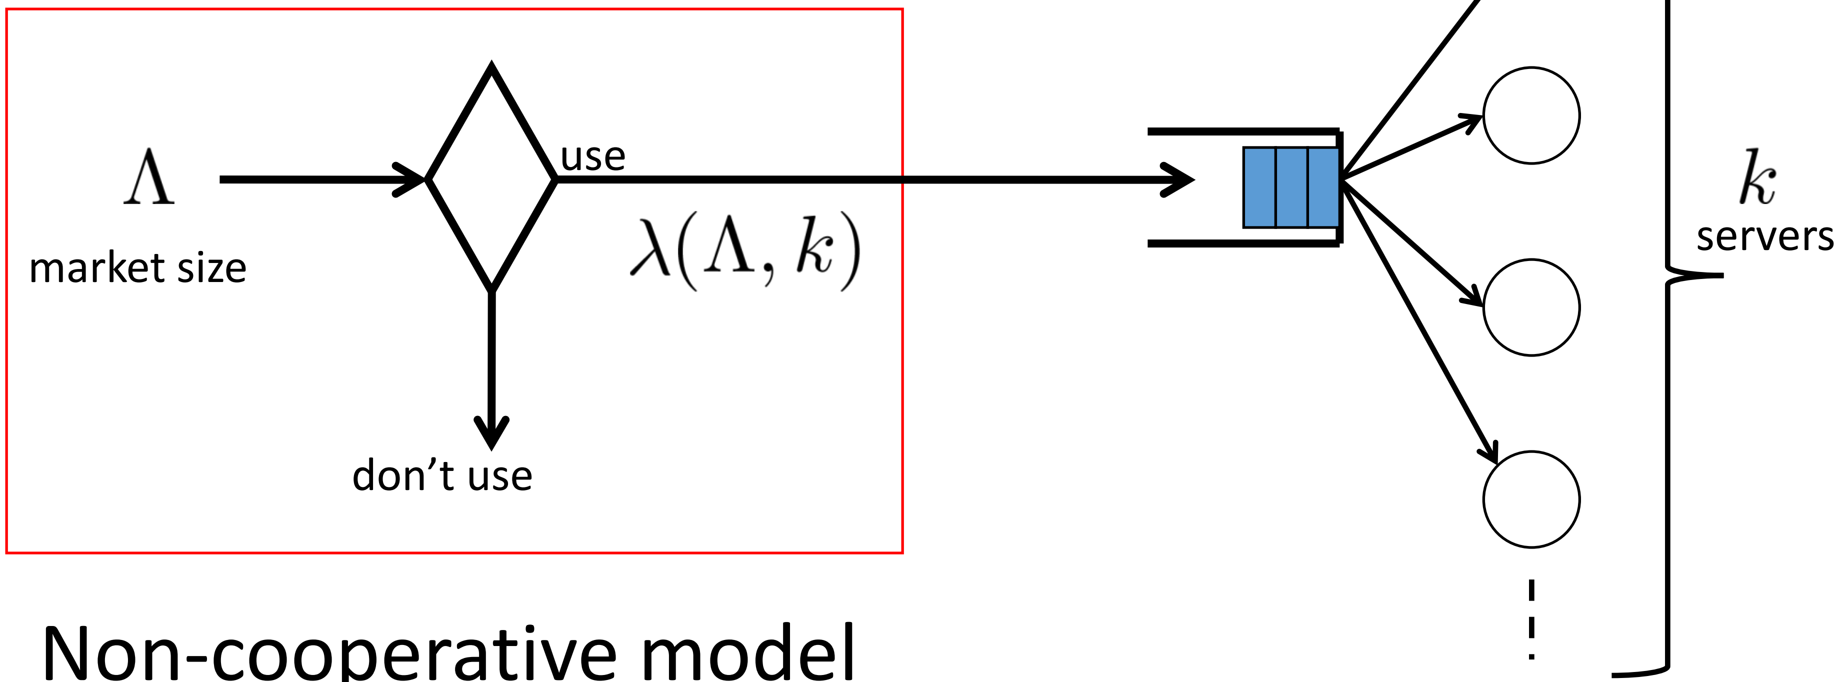

Non-cooperative model

$$\lambda(\Lambda, k) = \max\{\lambda \in [0, \Lambda] \mid V(\lambda) - f(\lambda, k) \geq 0\}$$

User utility

Congestion disutility

Wardrop equilibrium between users

Model (user base)

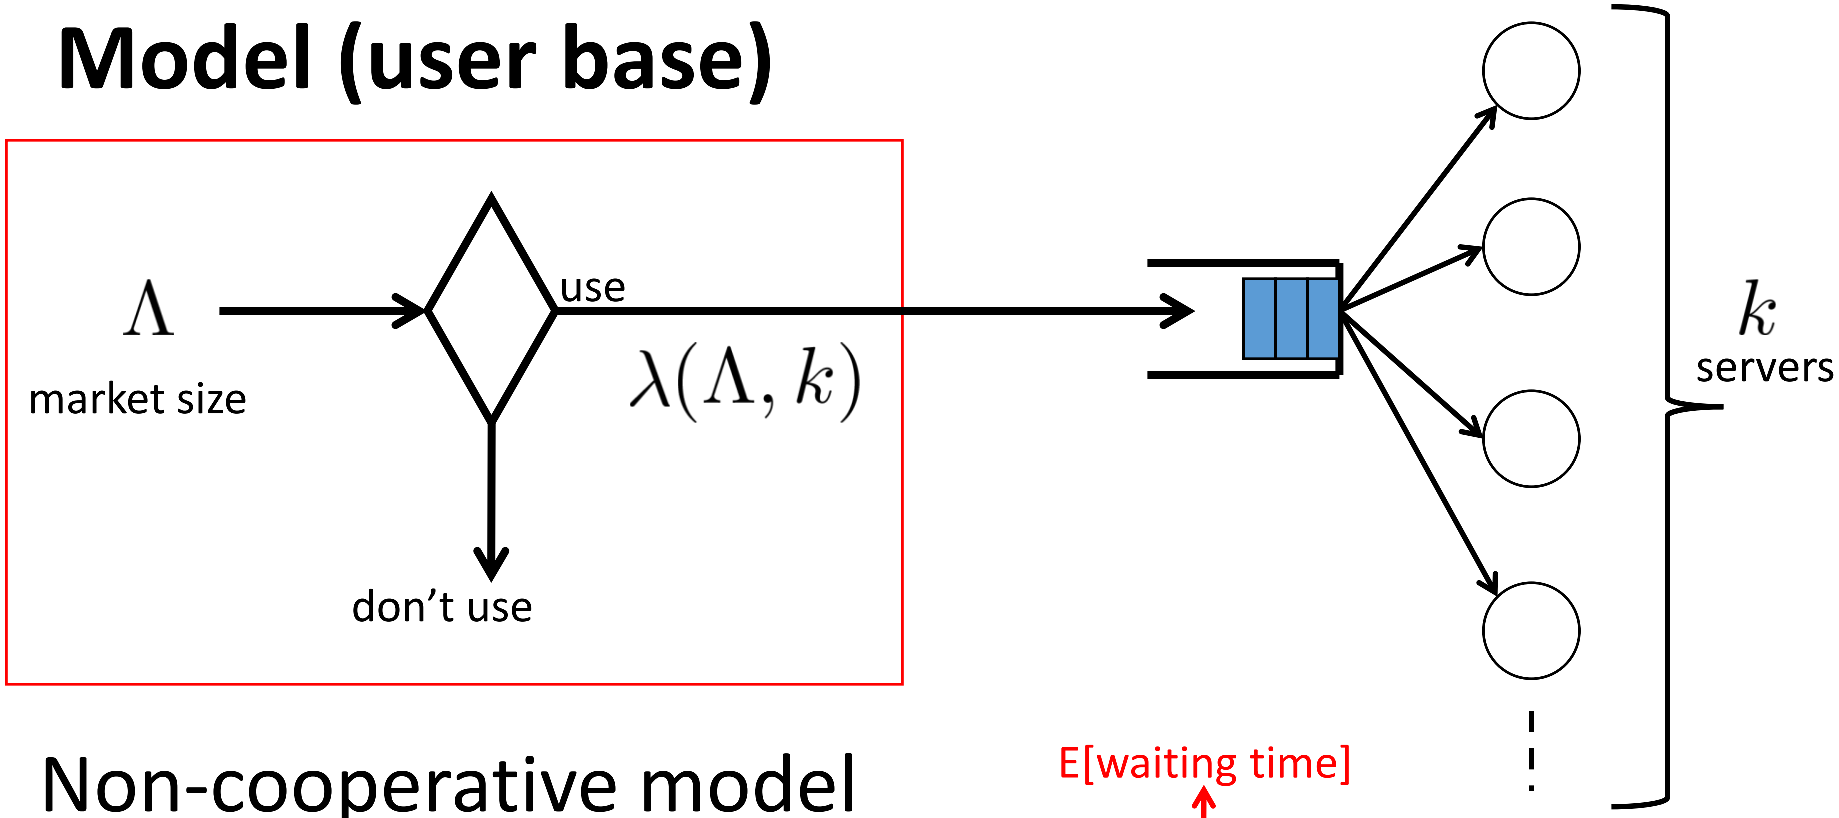

Non-cooperative model

E[waiting time]

$$\lambda(\Lambda, k) = \max\left\{\lambda \in [0, \Lambda] \mid V(\lambda) - f(\lambda, k) \geq 0\right\}$$

Model (user base)

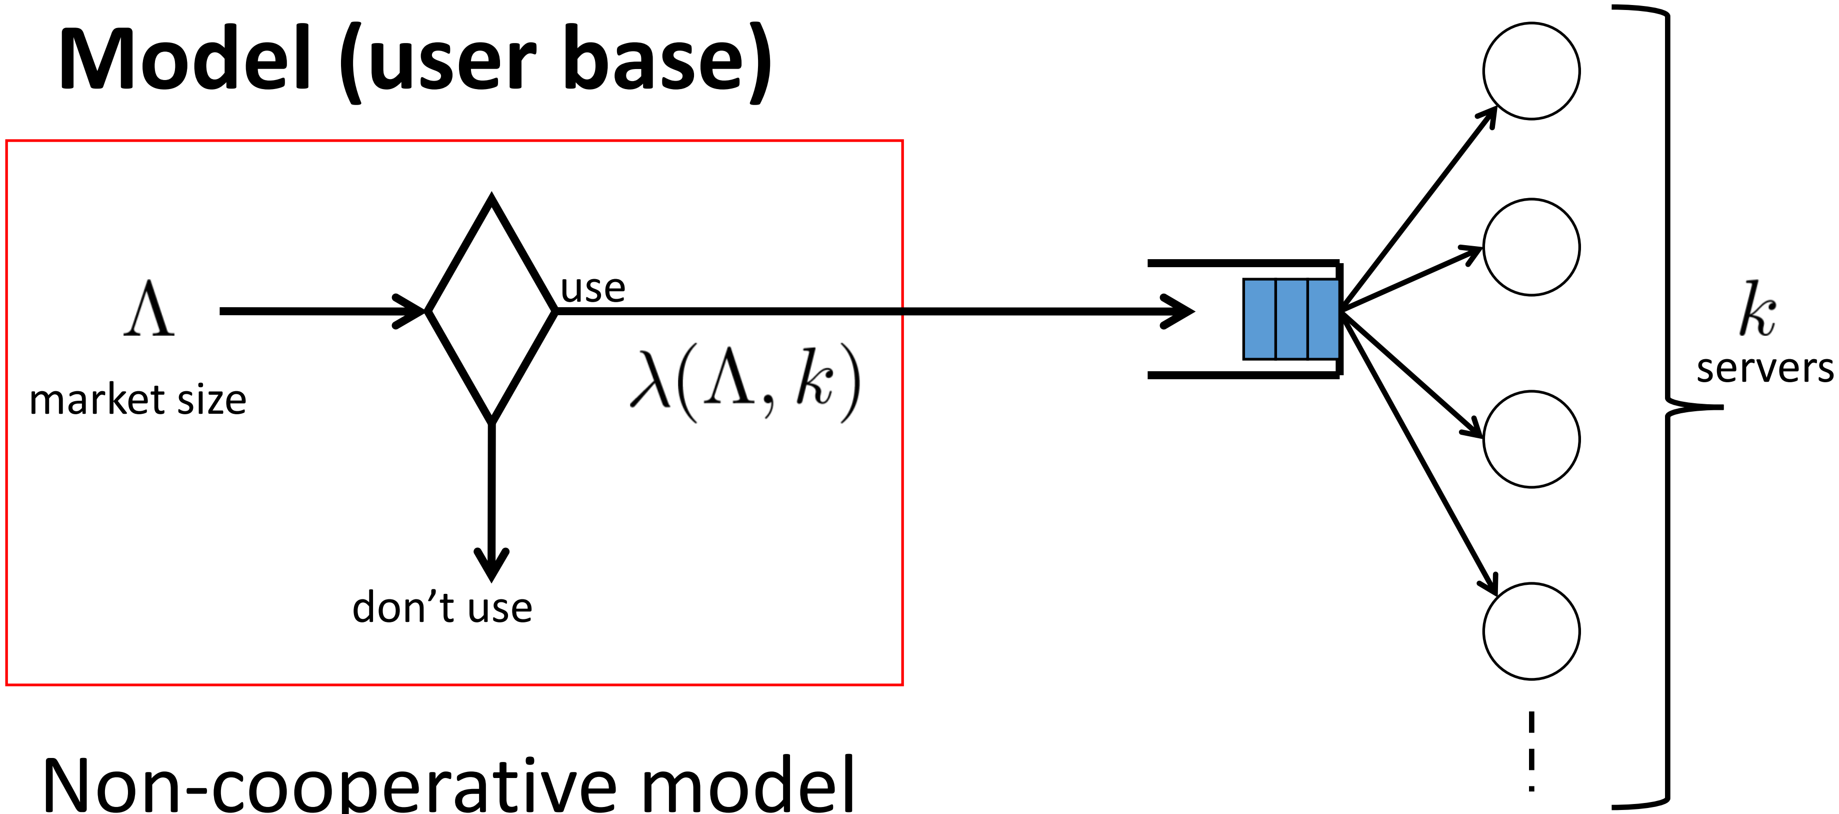

Non-cooperative model

$$\lambda(\Lambda, k) = \max\{\lambda \in [0, \Lambda] \mid V(\lambda) - f(\lambda, k) \geq 0\}$$

Determined by network effects

$$V(\lambda) = w\lambda^{\beta},\ w > 0,\ \beta \in [0, 1]$$

Model (user base)

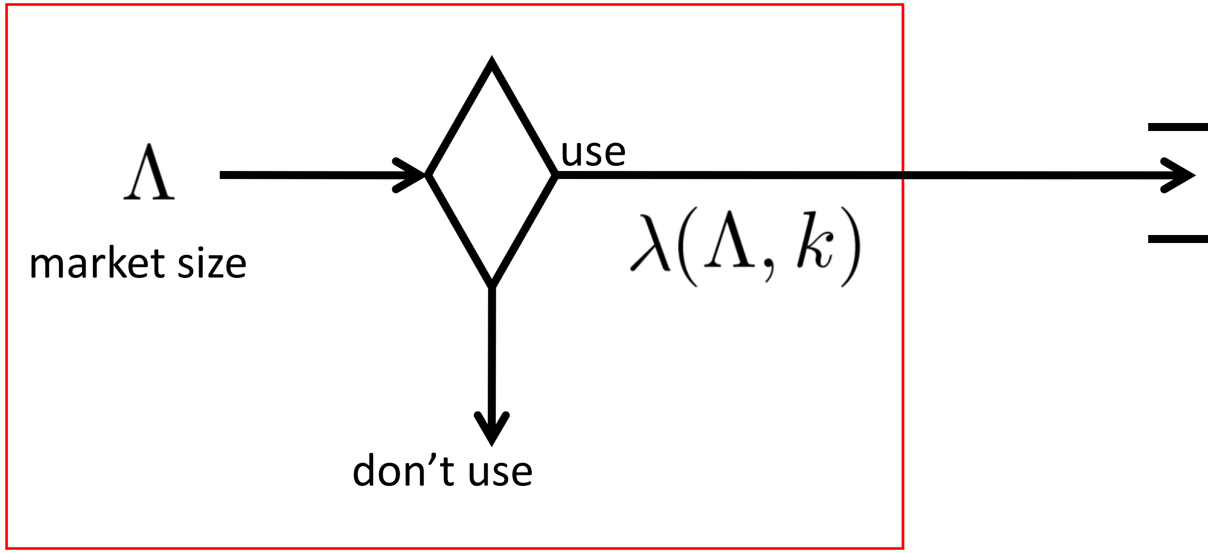

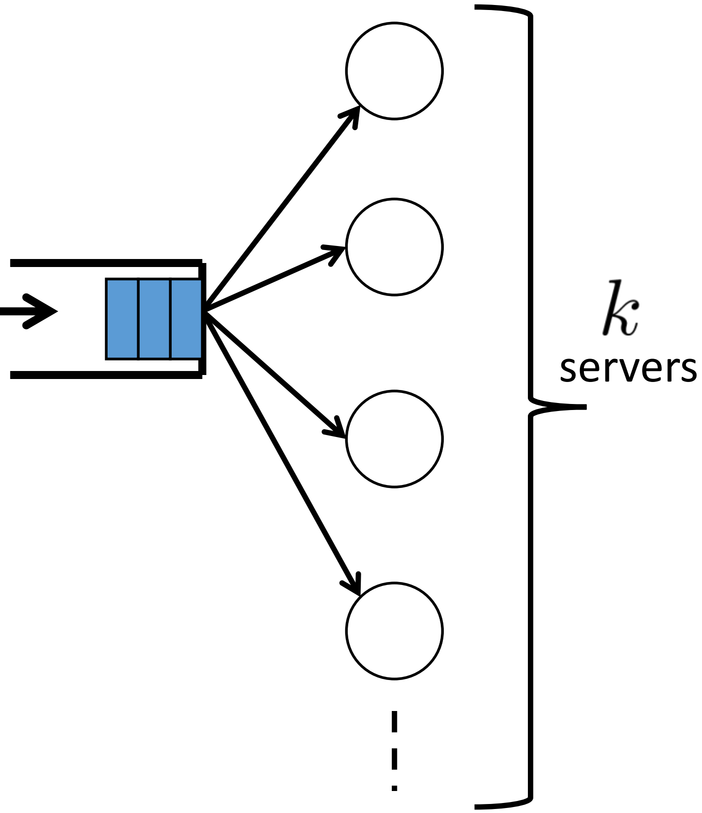

Non-cooperative model

$$\lambda(\Lambda, k) = \max\{\lambda \in [0, \Lambda] \mid V(\lambda) - f(\lambda, k) \geq 0\}$$

Cooperative model

$$\lambda(\Lambda, k) = \arg\max_{\lambda \in [0,\Lambda]} [\underbrace{\lambda V(\lambda)}_{\text{Aggregate utility}} - \underbrace{\lambda f(\lambda, k)}_{\text{Aggregate disutility}}]$$

Usage rate set to maximize global payoff

Model (service provider)

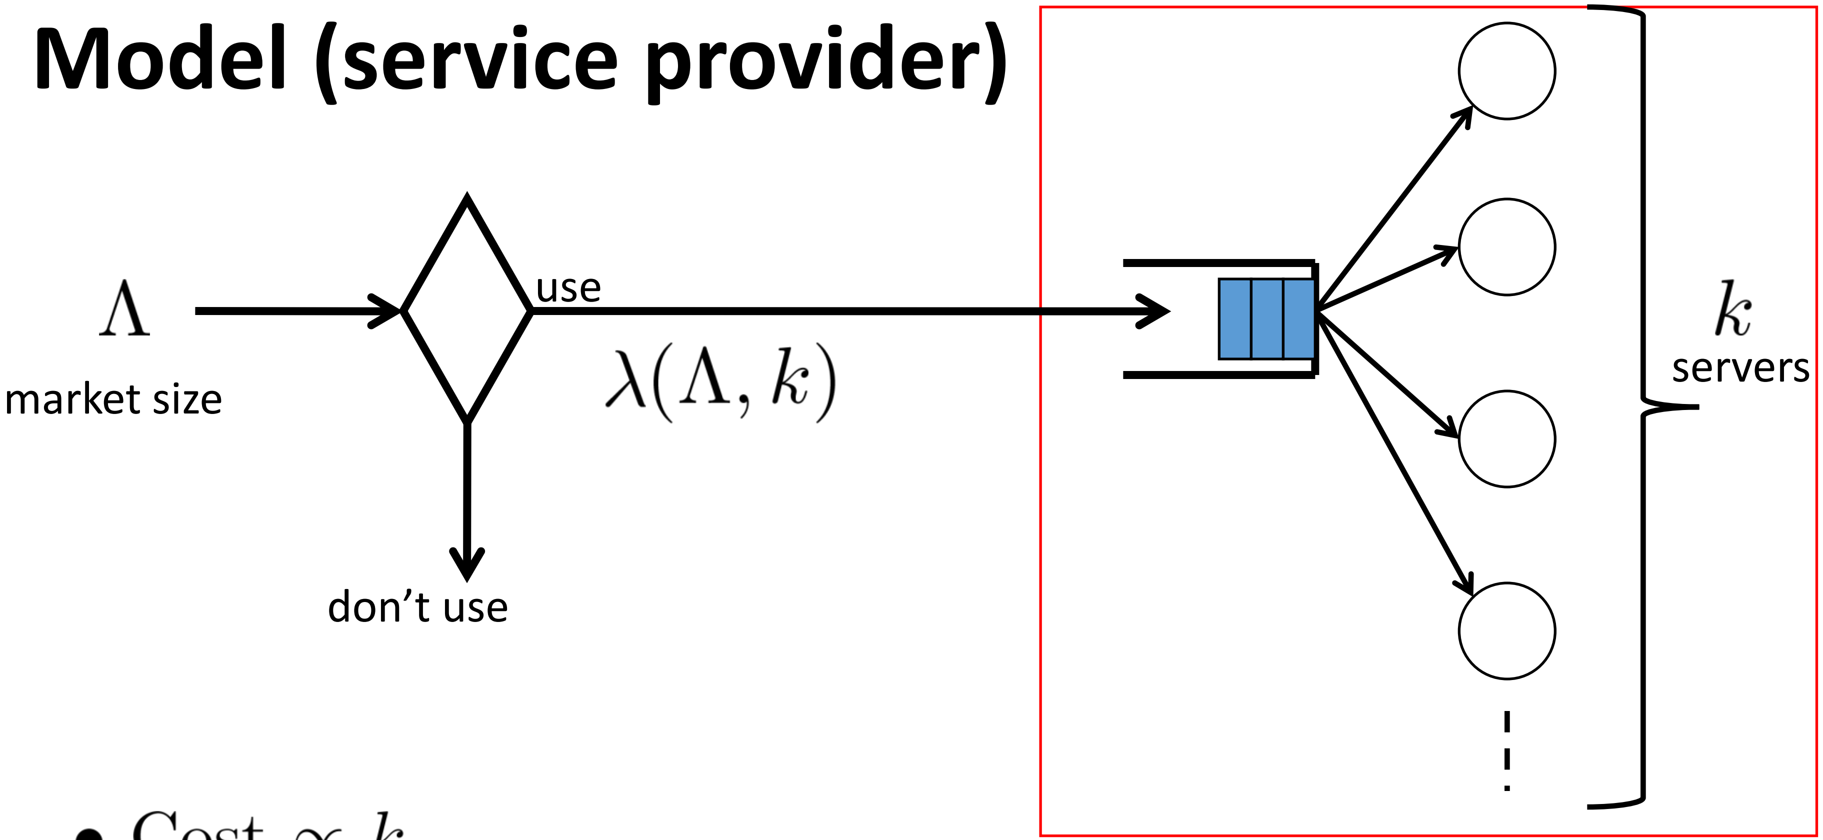

- Cost $\propto k$
- Revenue $\propto \lambda(\Lambda, k)$
- Profit maximizing strategy:

$$k^*_\Lambda = \arg\max_{k \geq 0} \left[b\lambda(\Lambda, k) - k \right]$$

Interesting case: $b > 1$

Model

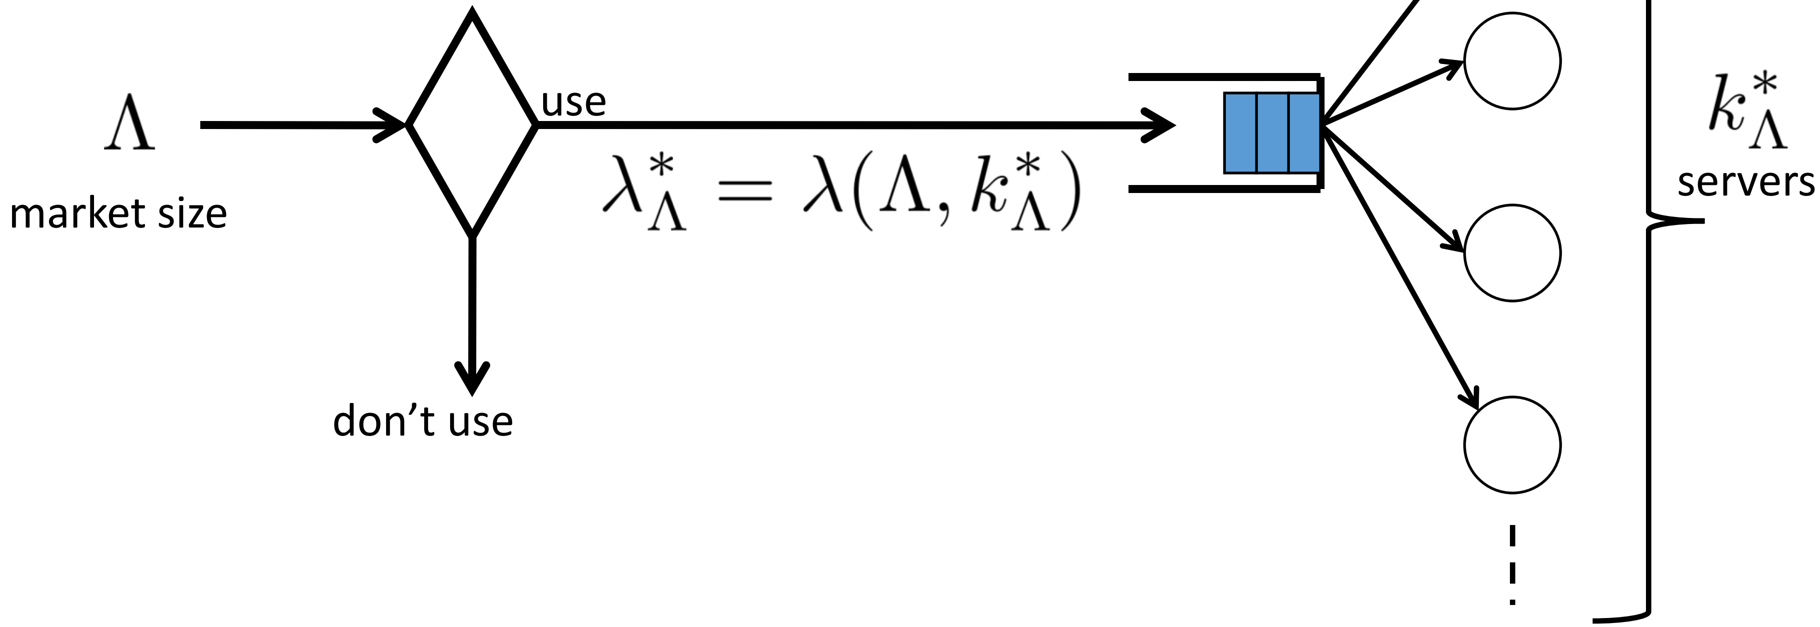

Want to understand:

- Capacity provisioning k_Λ^*
- $(\lambda_\Lambda^*, k_\Lambda^*)$
- Profit of the service provider

As a function of Λ

As $\Lambda \uparrow \infty$

(for large Λ)

Cooperative model

Theorem: For large enough Λ, $\lambda_\Lambda^* = \Lambda$. As $\Lambda \uparrow \infty$,

(a) if $\beta = 0$ (i.e., $V(\lambda) = w$),

$$k_\Lambda^* = \Lambda + \sqrt{\xi(w)\ \Lambda} + o(\sqrt{\Lambda})$$

$$\text{Profit} = \underbrace{(b-1)\Lambda}_{\text{Max. possible profit}} - \sqrt{\xi(w)\Lambda} - o(\sqrt{\Lambda})$$

Need 'square-root' spare servers
Halfin-Whitt regime

Cooperative model

Theorem: For large enough Λ, $\lambda^*_\Lambda = \Lambda$. As $\Lambda \uparrow \infty$,

(a) if $\beta = 0$ (i.e., $V(\lambda) = w$),

$$k^*_\Lambda = \Lambda + \sqrt{\xi(w)\ \Lambda} + o(\sqrt{\Lambda})$$

(b) if $\beta \in (0,1)$ (recall $V(\lambda) = w\lambda^\beta$),

$$k^*_\Lambda = \Lambda + \sqrt{\frac{\Lambda^{1-\beta}}{w(\beta+1)}} + o(\sqrt{\Lambda^{1-\beta}})$$

$$\text{Profit} = (b-1)\Lambda - \sqrt{\frac{\Lambda^{1-\beta}}{w(\beta+1)}} - o(\sqrt{\Lambda^{1-\beta}})$$

Less than 'square-root' spare servers

Cooperative model

Theorem: For large enough Λ, $\lambda^*_\Lambda = \Lambda$. As $\Lambda \uparrow \infty$,

(a) if $\beta = 0$ (i.e., $V(\lambda) = w$),

$$k^*_\Lambda = \Lambda + \sqrt{\xi(w)\ \Lambda} + o(\sqrt{\Lambda})$$

(b) if $\beta \in (0,1)$ (recall $V(\lambda) = w\lambda^\beta$),

$$k^*_\Lambda = \Lambda + \sqrt{\frac{\Lambda^{1-\beta}}{w(\beta+1)}} + o(\sqrt{\Lambda^{1-\beta}})$$

(c) if $\beta = 1$ (recall $V(\lambda) = w\lambda$),

$$k^*_\Lambda = \Lambda + O(1)$$

$$\text{Profit} = (b-1)\Lambda - O(1)$$

bounded spare servers

Cooperative model

Theorem: For large enough Λ, $\lambda^*_\Lambda = \Lambda$. As $\Lambda \uparrow \infty$,

(a) if $\beta = 0$ (i.e., $V(\lambda) = w$),

$$k^*_\Lambda = \Lambda + \sqrt{\xi(w)\ \Lambda} + o(\sqrt{\Lambda})$$

(b) if $\beta \in (0,1)$ (recall $V(\lambda) = w\lambda^\beta$),

$$k^*_\Lambda = \Lambda + \sqrt{\frac{\Lambda^{1-\beta}}{w(\beta+1)}} + o(\sqrt{\Lambda^{1-\beta}})$$

(c) if $\beta = 1$ (recall $V(\lambda) = w\lambda$),

$$k^*_\Lambda = \Lambda + O(1)$$

Stronger network effects $\Longrightarrow$ less service capacity, increased profit

Non-cooperative model

Theorem: For large enough Λ, $\lambda^*_\Lambda = \Lambda$. As $\Lambda \uparrow \infty$,

$$k^*_\Lambda = \Lambda + O(1)$$

Tragedy of the commons: Bounded spare servers irrespective of network effects

Impact of network effects heavily diminished by anarchy in user base

Proof idea

Original system

$$\lambda(\Lambda, k) = \arg\max_{\lambda \in [0,\Lambda]} \left[\lambda V(\lambda) - \lambda \mathbb{E}[W]\right]$$

$$k^*_\Lambda = \arg\max_{k \geq 0} \left[b\lambda(\Lambda, k) - k\right]$$

Unconstrained system

$$\boldsymbol{\lambda}(k) = \arg\max_{\lambda \geq 0} \left[U(\lambda) - \lambda \mathbb{E}[W]\right]$$

Want to understand

$(\lambda^*_\Lambda, k^*_\Lambda)$ ⟷ $(\boldsymbol{\lambda}(k), k)$

Characterize

$$k^*_\Lambda \approx \boldsymbol{\lambda}^{-1}(\Lambda)$$

$$\lambda^*_\Lambda \approx \Lambda$$

Case of competing providers

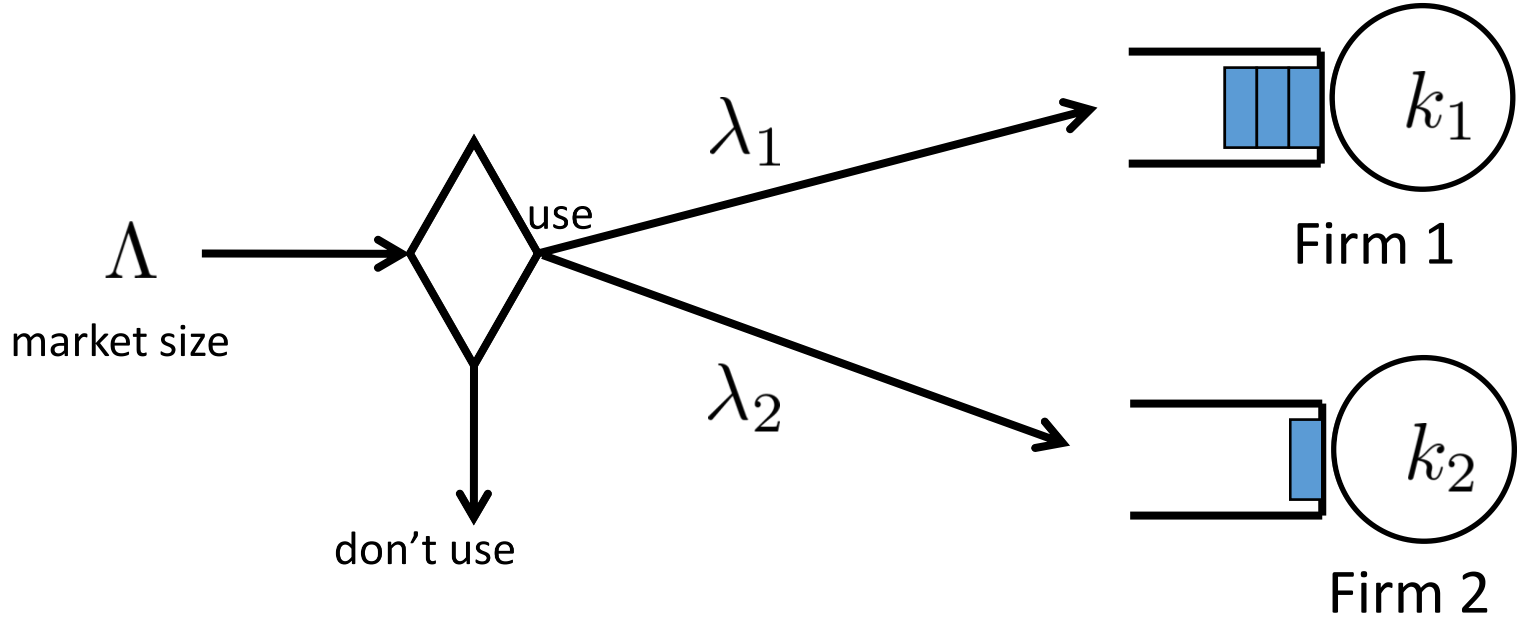

Each firm operates an M/M/1 system

Focus on non-cooperative user behavior

Two cases depending on nature of network effects:

1. Industrywide network effects: Utility from Firm $i = V(\lambda_1 + \lambda_2)$
2. Firm-specific network effects: Utility from Firm $i = V_i(\lambda_i)$

Industry-wide network effects

Model (user base)

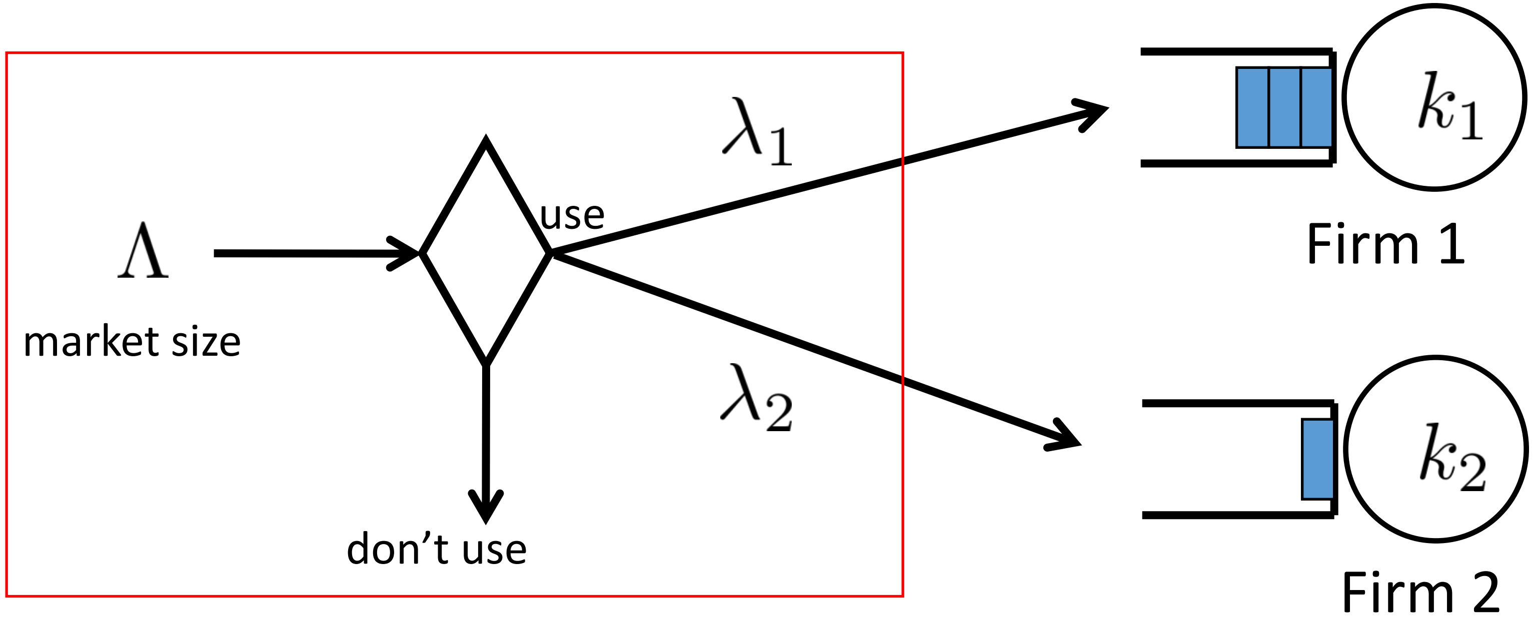

Given $k = (k_1, k_2)$, $\lambda < k_1 + k_2$, Wardrop split $(\hat{\lambda}_1(\lambda, k), \hat{\lambda}_2(\lambda, k))$ solves

$$\max_{\lambda_1, \lambda_2} \sum_{i=1}^{2} \int_0^{\lambda_i} [V(\lambda) - f(x, k_i)] dx$$

$$\text{subject to } \sum_{i=1}^{2} \lambda_i = \lambda$$

Model (user base)

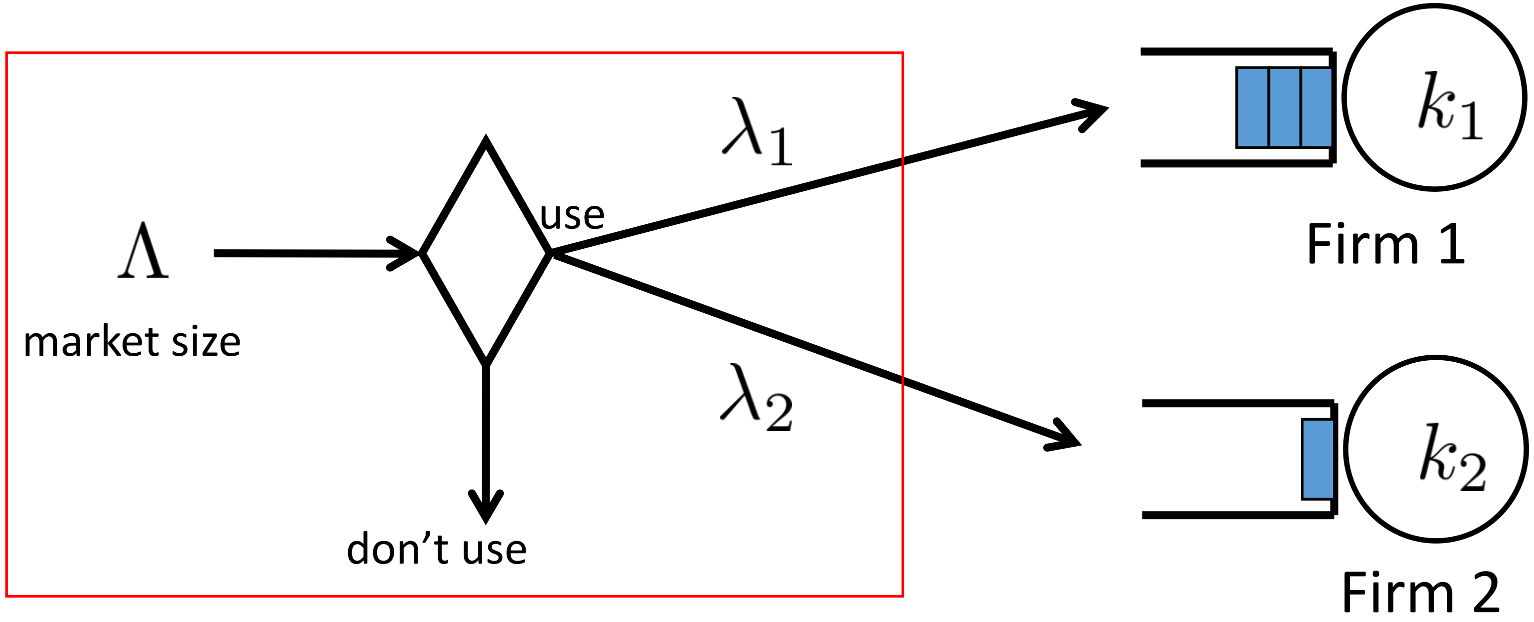

$$\lambda(k) = \max\left\{\lambda \in [0, \Lambda] \cap [0, k_1 + k_2) \mid V(\lambda) - f(\hat{\lambda}_1(\lambda, k), k_1) \geq 0\right\}$$

$$\lambda_1(k) = \hat{\lambda}_1(\lambda(k), k), \quad \lambda_2(k) = \hat{\lambda}_2(\lambda(k), k)$$

Model (firms)

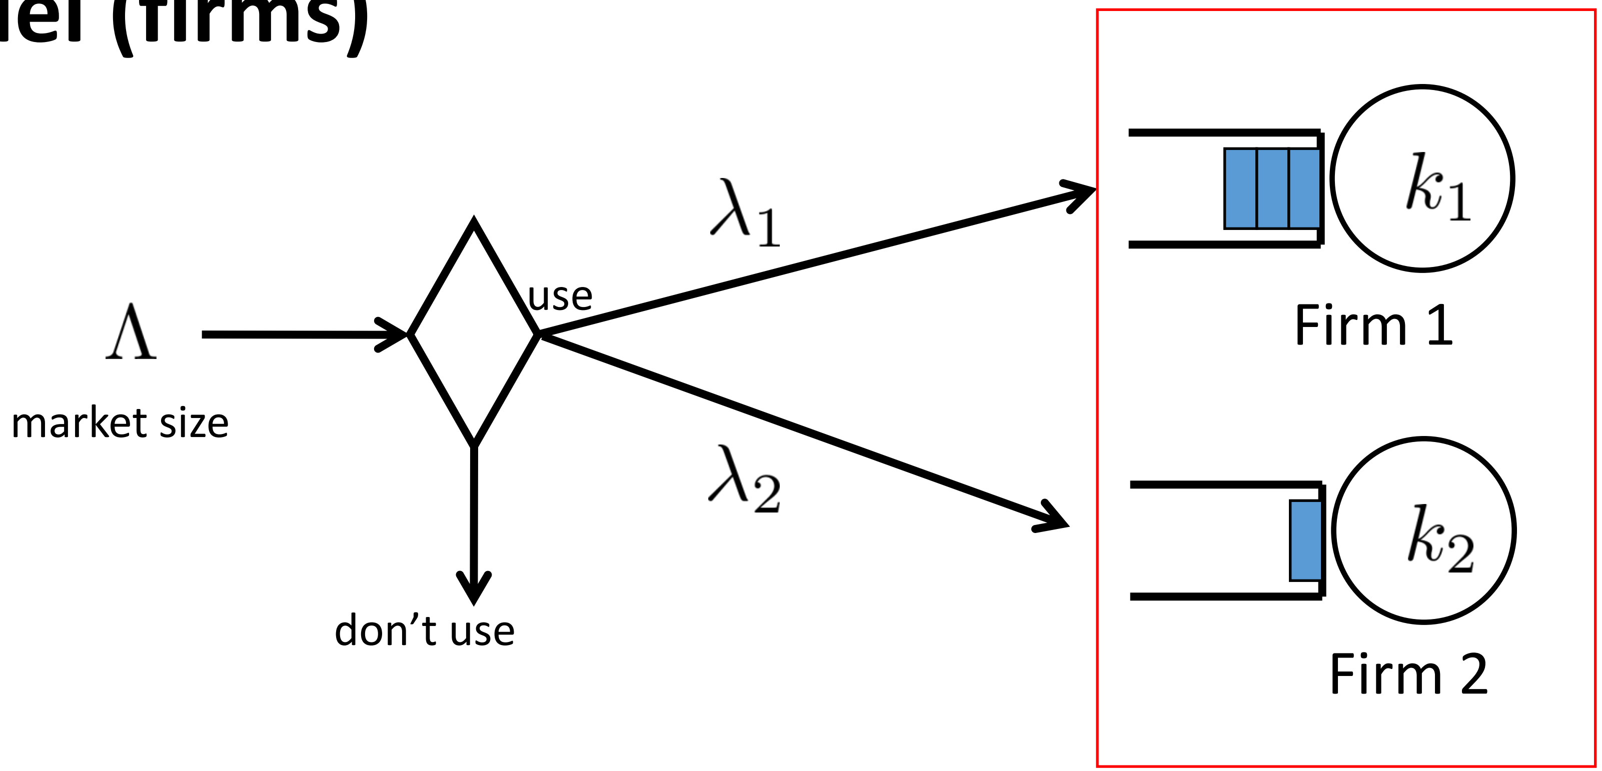

$$\text{Firm } i \text{ profit} = b_i \lambda_i(k) - k_i$$

We look for Nash equilibria

Theorem: If $\boldsymbol{b_1, b_2 \in (1, 2]}$, then a continuum of equilibria exist, including monopoly configurations. Any equilibrium is of one of the following forms.

1. Monopoly for Firm 1: $\lambda_1 = \Lambda$, $C_1 = \Lambda + \frac{1}{V(\Lambda)}$, $\lambda_2 = C_2 = 0$
2. Monopoly for Firm 2: $\lambda_2 = \Lambda$, $C_2 = \Lambda + \frac{1}{V(\Lambda)}$, $\lambda_1 = C_1 = 0$
3. Firms 1 and 2 share the market such that $\lambda_1 + \lambda_2 = \Lambda$, and

$$\lambda_i \geq \frac{1}{(b_i - 1)V(\Lambda)},$$
$$C_i = \lambda_i + \frac{1}{V(\Lambda)}$$

When network effects are industry-wide, multiple firms can share the market.

Competition between firms does not help the user base

Theorem (contd.): If $\boldsymbol{b_1, b_2 > 2}$, there there is no Nash equilibrium.
If $\boldsymbol{b_1 > 2, b_2 \leq 2}$, then only equilibrium is monopoly of Firm 1:

$$\lambda_1 = \Lambda, k_1 = \Lambda + \frac{1}{V(\Lambda)}, \lambda_2 = k_2 = 0$$

Firm-specific network effects

Model (user base)

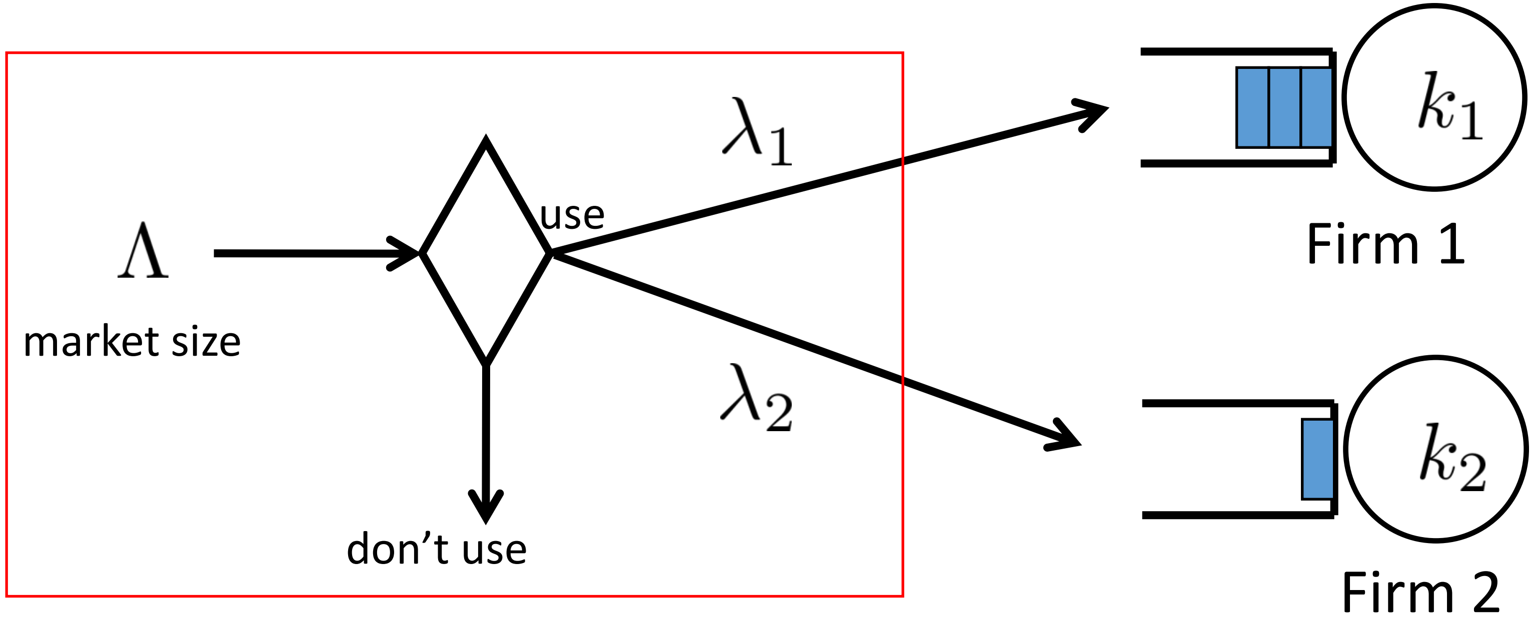

Here, $V_i(\lambda_i) = w_i$

WLOG, assume $w_1 > w_2$

Model (user base)

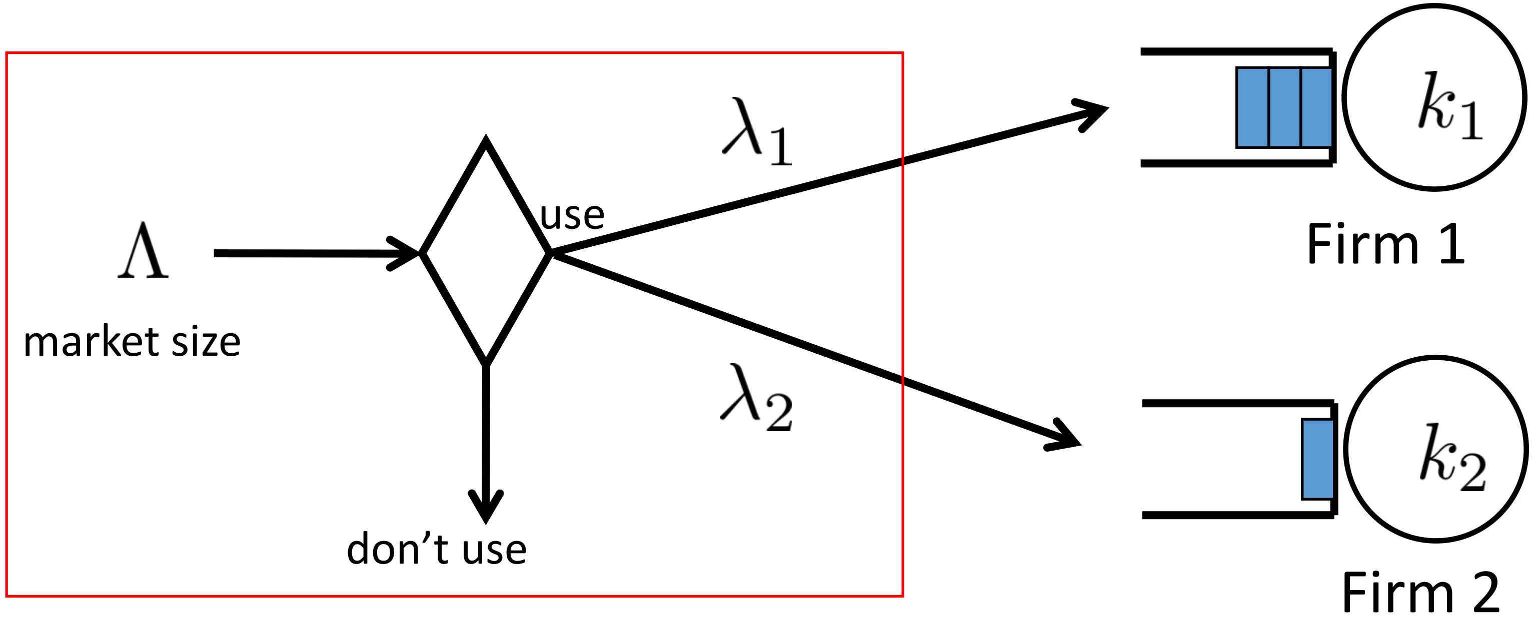

Given $k = (k_1, k_2)$, $\lambda < k_1 + k_2$, Wardrop split $(\hat{\lambda}_1(\lambda, k), \hat{\lambda}_2(\lambda, k))$ solves

$$\max_{\lambda_1, \lambda_2} \sum_{i=1}^{2} \int_0^{\lambda_i} [w_i - f(x, k_i)] dx$$

$$\text{subject to } \sum_{i=1}^{2} \lambda_i = \lambda$$

Model (user base)

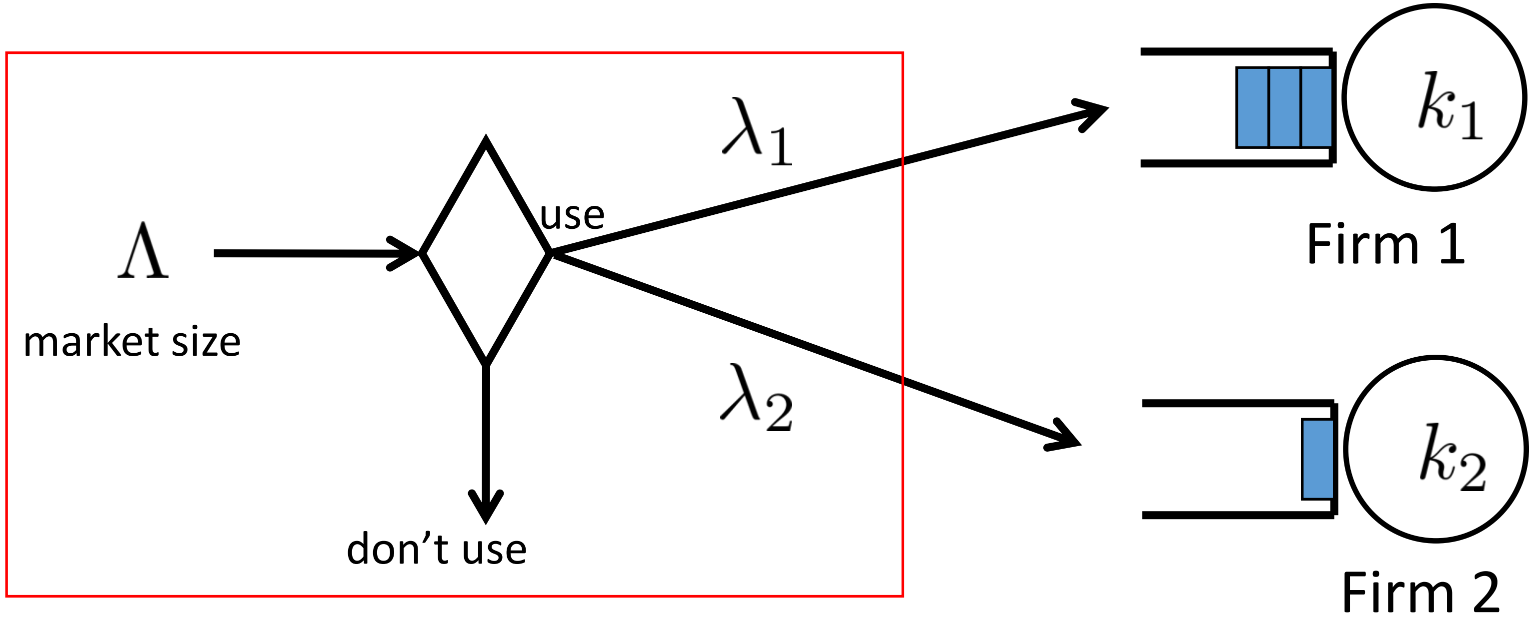

$$\lambda(k) = \max\left\{\lambda \in [0,\Lambda] \cap [0, k_1 + k_2) \mid w_1 - f(\hat{\lambda}_1(\lambda, k), k_1) \geq 0\right\}$$

$$\lambda_1(k) = \hat{\lambda}_1(\lambda(k), k), \quad \lambda_2(k) = \hat{\lambda}_2(\lambda(k), k)$$

Model (firms)

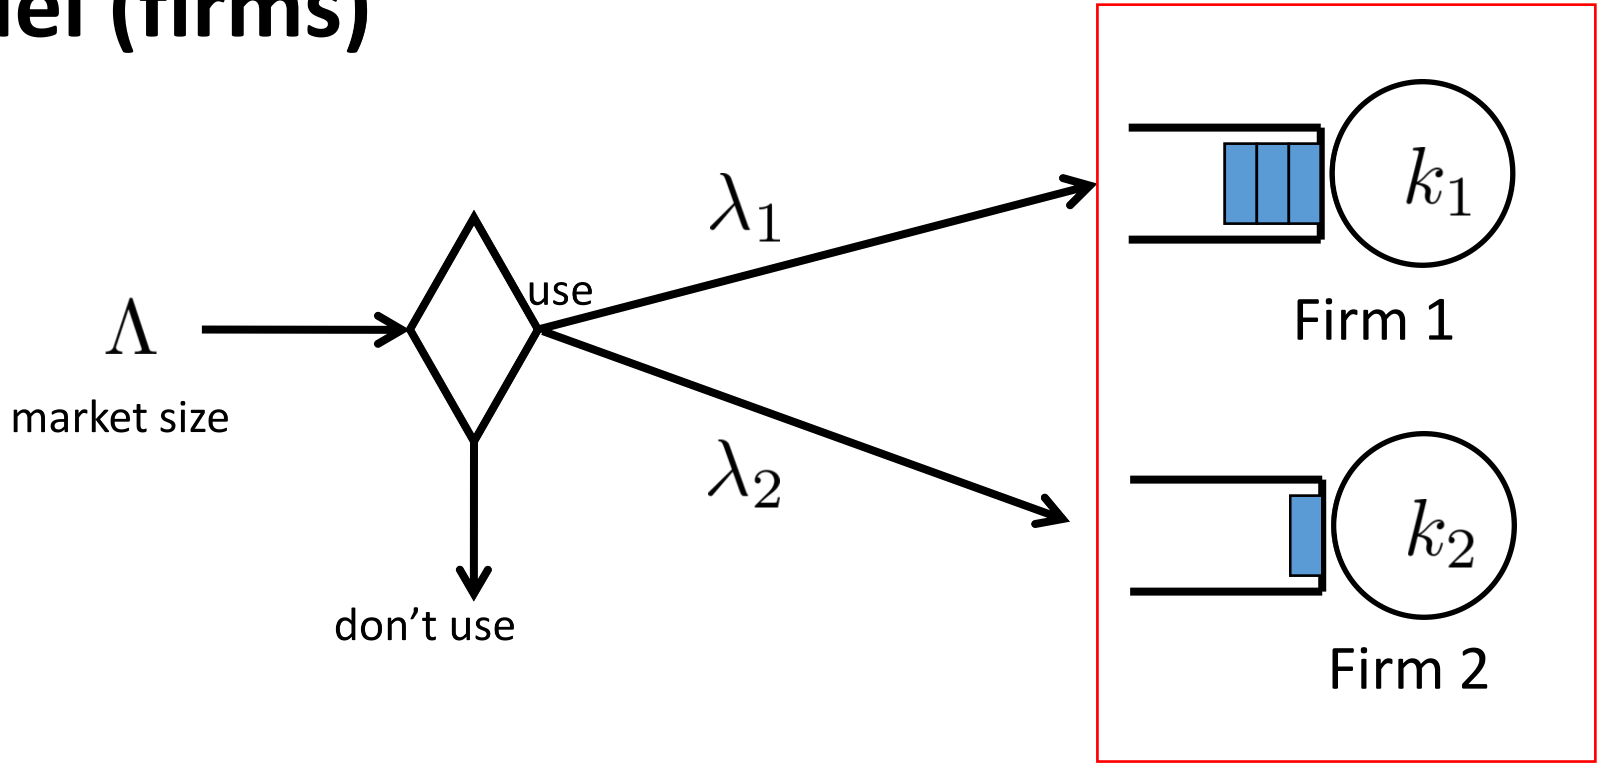

$$\text{Firm } i \text{ profit} = b_i \lambda_i(k) - k_i$$

As before, we look for Nash equilibria

Theorem: Let $w_1 > w_2 > 0$. For large enough Λ, any equilibrium must satisfy

$$\lambda_1 \geq \Lambda - \frac{1}{b_1 - 1}\left(\frac{b_1 w_2}{w_1(w_1 - w_2)} + \frac{1}{w_1}\right).$$

Near monopoly for Firm 1

Proof idea: For any capacity k_2, Firm 1 can provision as much capacity as the single firm case and attract most of the user base

Summary

- Industry-wide network effects => firms can share market
  Firm specific network effects => near monopolies
- Competition between firms does not help the user base

Competitive Provisioning of Online Services

Jayakrishnan Nair

EE, IITB

Joint work with Vijay Subramanian, Adam Wierman & Bert Zwart