Asymptotically Optimal Control of a Centralized Dynamic Matching Market with General Utilities

Jose H. Blanchet, Virag Shah, Linjia Wu
Management Science and Engineering Department
Stanford University

Martin I. Reiman
Industrial Engineering and Operations Research Department
Columbia University

Lawrence M. Wein
Graduate School of Business
Stanford University

Introduction

- Taxonomy
  - Static vs. Dynamic
  - Centralized vs. Decentralized
  - Binary utilities vs. General utilities

Introduction

- Taxonomy
  - Static vs. Dynamic
  - Centralized vs. Decentralized
  - Binary utilities vs. General utilities

- Most dynamic MM models (except, e.g., Unver 2010 and Hu & Zhou 2016) assume dichotomous match outcomes

Introduction

- Taxonomy
  - Static vs. Dynamic
  - Centralized vs. Decentralized
  - Binary utilities vs. General utilities

- Most dynamic MM models (except Unver 2010 and Hu & Zhou 2016) assume dichotomous match outcomes

- Closest related work:
  - Hu & Zhou (2016): Structural results for multiclass model
  - Liu, Gong & Kulkarni (2015), Busic & Meyn (2016) and Buke & Chen (2017): fluid and diffusion limits for simpler problems
  - Mertikopoulos et al. (2020): use $\pi^2/6$ result (Mezard-Parisi-Aldous) to study batch-and-match policies to minimize exponential mismatch plus waiting costs

Symmetric Model

- Buyers and sellers each arrive according to independent Poisson processes with rate λ
- Each buyer and seller independently abandons the market after an exponential amount of time with rate η

Symmetric Model

- Buyers and Sellers each arrive according to independent Poisson processes with rate λ
- Each buyer and seller independently abandons the market after an exponential amount of time with rate η
- Utility of a match between any buyer and any seller is $V \sim F(v)$
- The goal is to maximize the long run average expected utility

Symmetric Model

- Buyers and Sellers each arrive according to independent Poisson processes with rate λ
- Each buyer and seller independently abandons the market after an exponential amount of time with rate η
- Utility of a match between any buyer and any seller is $V \sim F(v)$
- The goal is to maximize the long run average expected utility
- The key tradeoff: make a match now or wait for a better match later?

A Restricted Class of Policies

- Buyer arrives at time t to find $S(t)$ sellers, and system manager observes $V_1, \ldots, V_{S(t)}$
- Seller arrives at time t to find $B(t)$ buyers, and system manager observes $V_1, \ldots, V_{B(t)}$

A Restricted Class of Policies

- Buyer arrives at time t to find $S(t)$ sellers, and system manager observes $V_1, \ldots, V_{S(t)}$
- Seller arrives at time t to find $B(t)$ buyers, and system manager observes $V_1, \ldots, V_{B(t)}$
- Matches can be made only at an arrival epoch, and must involve the arriving agent

Large-Market Scaling

Even under a simple control policy, this model gives rise to a two-dimensional continuous time Markov chain (CTMC), $(B(t), S(t))$, which is difficult to deal with

So we consider asymptotics

Large-Market Scaling

Even under a simple control policy, this model gives rise to a two-dimensional continuous time Markov chain (CTMC), $(B(t), S(t))$, which is difficult to deal with

So we consider asymptotics

In n^{th} system (as $n \to \infty$)

- Arrival rates $= n\lambda$
- Abandonment rates $= \eta$ (unscaled)
- Matching values $= V$ (unscaled)
- CTMC state $= (B_n(t), S_n(t))$
- $O(n)$ agents in system if no matches

Extreme Value Theory

- $M_n = \max\{V_1, \ldots, V_n\}$
- For Types $j =$ I, II, III

$$P\left(\frac{M_n - b_n}{a_n} \leqslant x\right) \to G_j(x) \quad as \quad n \to \infty,$$

Extreme Value Theory

- $M_n = \max\{V_1, \ldots, V_n\}$
- For Types $j =$ I, II, III

$$P\left(\frac{M_n - b_n}{a_n} \leqslant x\right) \to G_j(x) \quad as \quad n \to \infty,$$

- Type I = Gumbel (e.g., exponential, normal, gamma, lognormal)
- Type II = Frechet (e.g., Pareto)
- Type III = Reverse Weibull (e.g., uniform, beta)
- Interested primarily in $E[M_n] \sim a_n \mu_j + b_n$ for $j =$ I, II, III

Population-based Threshold Policy With Threshold z_n

- If buyer arrives and finds $S_n(t) > z_n$ sellers, then he matches to seller with highest matching utility

Population-based Threshold Policy With Threshold z_n

- If buyer arrives and finds $S_n(t) > z_n$ sellers, then he matches to seller with highest matching utility
- If buyer arrives and finds $S_n(t) \leqslant z_n$ sellers, then he waits in market

Population-based Threshold Policy With Threshold z_n

- If buyer arrives and finds $S_n(t) > z_n$ sellers, then he matches to seller with highest matching utility
- If buyer arrives and finds $S_n(t) \leqslant z_n$ sellers, then he waits in market
- Greedy policy corresponds to $z_n = 0$

Population-based Threshold Policy With Threshold z_n

- If buyer arrives and finds $S_n(t) > z_n$ sellers, then he matches to seller with highest matching utility
- If buyer arrives and finds $S_n(t) \leqslant z_n$ sellers, then he waits in market
- Greedy policy corresponds to $z_n = 0$
- Similar match/wait decision when seller arrives to find $B_n(t)$ buyers

The Utility Rate U_n

- $U_n =$ arrival rate $\times$ P(agent is matched) $\times$ E[utility per match]
- Arrival rate $= n\lambda$
- P(agent is matched) is derived from queueing asymptotics
- E[utility per match] is derived from extreme value theory

The Utility Rate U_n

- U_n = arrival rate $\times$ P(agent is matched) $\times$ E[utility per match]
- Arrival rate $= n\lambda$
- P(agent is matched) is derived from queueing asymptotics
- E[utility per match] is derived from extreme value theory
- Key Lemma: $E[M_{B_n}] = E[M_{E[B_n]}]$ in fluid limit
- Upper bound
    - P(agent is matched) $= 1$
    - E[utility/match]: assume no matching in $(B_n(t), S_n(t)) \Rightarrow$ arriving buyers see Poi$\left(\frac{n\lambda}{\eta}\right)$ sellers, and matches to best one
    - $U_n^u \sim n\lambda E[M_{\frac{n\lambda}{\eta}}]$

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
exponential (ν)	$U_n^u \sim \frac{\lambda}{\nu} n \ln n$		

- $U_n^u \sim n\lambda E[M_{\frac{n\lambda}{\eta}}]$
- $E[M_n] \sim \frac{\gamma + \ln n}{\nu}$ where $\gamma = 0.5772$

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
exponential (ν)	$U_n^u \sim \frac{\lambda}{\nu} n \ln n$	$U_n^g \sim \frac{\lambda}{2\nu} n \ln n$	

- Steady-state distribution of $\frac{B_n(t) - S_n(t)}{\sqrt{n}} \to N\left(0, \frac{\lambda}{\eta}\right)$ (Liu et al. 2015)
- $\Pr(\text{abandon}) = \frac{abandonment\ \ rate}{arrival\ \ rate} = \frac{O(\sqrt{n})}{O(n)} \to 0$
- $U_n^g \sim n\lambda E[M_{\frac{\lambda}{\eta}\sqrt{\frac{2}{\pi}}\sqrt{n}}]$

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
exponential (ν)	$U_n^u \sim \frac{\lambda}{\nu} n \ln n$	$U_n^g \sim \frac{\lambda}{2\nu} n \ln n$	$z_n^* = \frac{n}{\ln n}$ is asymptotically optimal

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
exponential (ν)	$U_n^u \sim \frac{\lambda}{\nu} n \ln n$	$U_n^g \sim \frac{\lambda}{2\nu} n \ln n$	$z_n^* = \frac{n}{\ln n}$ is asymptotically optimal

More general framework: $E[M_t]$ is regularly varying at ∞ with index $\alpha \in [0,1)$; i.e., $\lim_{t\to\infty} \frac{m(tx)}{m(t)} = x^\alpha$

Theorem: Assume that $\alpha = 0$ and let $l(n) > 0$ be any slowly varying function at ∞ such that $l(n) \nearrow \infty$ as $n \to \infty$. Then a threshold policy with $z_n = \frac{n}{l(n)}$ is asymptotically optimal.

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
Pareto shape $\beta > 1$	$U_n^u = O(n^{1+\frac{1}{\beta}})$	$U_n^g = O(n^{1+\frac{1}{2\beta}})$	

- Matching Utilities are Pareto(1,2), $F(v) = 1 - v^{-2}, v \geqslant 1$ $(\beta = 2)$
- $E[M_n] \sim \sqrt{n\pi}$
- $U_n^u \sim \frac{\sqrt{\pi}\lambda^{3/2}}{\sqrt{\eta}} n^{3/2}$ upper bound
- $U_n^g \sim (2\pi)^{1/4} \frac{\lambda^{3/2}}{\sqrt{\eta}} n^{5/4}$ greedy policy
- Threshold policy is much better than greedy policy $(n^{5/4}$ vs. $n^{3/2})$

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
Pareto shape $\beta > 1$	$U_n^u = O(n^{1+\frac{1}{\beta}})$	$U_n^g = O(n^{1+\frac{1}{2\beta}})$	$z_n^* = \frac{\lambda}{\eta(1+\beta)} n$ $U_n^t(z_n^*) = O(n^{1+\frac{1}{\beta}})$ does not converge to upper bound

- Matching Utilities are Pareto(1,2), $F(v) = 1 - v^{-2}, v \geqslant 1$ $(\beta = 2)$
- $E[M_n] \sim \sqrt{n\pi}$
- $\max_z \; U_n^t(nz) = \max_z \; n\lambda\left(1 - \frac{z\eta}{\lambda}\right)\sqrt{nz\pi}$ threshold policy

  $\Rightarrow z^* = \frac{\lambda}{3\eta}$ simple optimal threshold
- $U_n^t\left(\frac{\lambda n}{3\eta}\right) \sim \frac{2\sqrt{\pi}}{3\sqrt{3}} \frac{\lambda^{3/2}}{\sqrt{\eta}} n^{3/2}$ threshold policy
- Upper bound is loose (by factor $\frac{2}{3\sqrt{3}} = 0.385$)

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
Pareto shape $\beta > 1$	$U_n^u = O(n^{1+\frac{1}{\beta}})$	$U_n^g = O(n^{1+\frac{1}{2\beta}})$	$z_n^* = \frac{\lambda}{\eta(1+\beta)} n$ $U_n^t(z_n^*) = O(n^{1+\frac{1}{\beta}})$ does not converge to upper bound

More general framework: $E[M_t]$ is regularly varying at ∞ with index $\alpha \in [0, 1)$; i.e., $\lim_{t\to\infty} \frac{m(tx)}{m(t)} = x^\alpha$

Theorem: Assume that $\alpha \in (0, 1)$. Then a threshold policy of the form $z_n = z_* n$ with $z_* = \frac{\lambda\alpha}{\eta(1+\alpha)}$ is asymptotically optimal within the class of population-based threshold policies

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
uniform$[a, b]$	$U_n^u \sim \lambda bn$	asymptotically optimal	$z_n^* = 0$ is asymptotically optimal

- Matching Utilities are U[0,1]
- $E[M_n] \sim 1 - \frac{1}{n}$
- $U_n^u \sim n\lambda\left(1 - \frac{1}{\frac{\lambda}{\eta} n}\right)$ upper bound
  $\sim \lambda n$
- $U_n^g \sim n\lambda\left(1 - \frac{1}{\frac{\lambda}{\eta}\sqrt{\frac{2}{\pi}}\sqrt{n}}\right)$ greedy policy
  $\sim \lambda n$

Summary of Results for Three Canonical Examples

Matching Utility Distribution	Utility Rate: Upper Bound	Utility Rate: Greedy Policy	Utility Rate: Threshold Policy
uniform$[a, b]$	$U_n^u \sim \lambda b n$	asymptotically optimal	$z_n^* = 0$ is asymptotically optimal
exponential (ν)	$U_n^u \sim \frac{\lambda}{\nu} n \ln n$	$U_n^g \sim \frac{\lambda}{2\nu} n \ln n$	$z_n^* = \frac{n}{\ln n}$ is asymptotically optimal
Pareto shape $\beta > 1$	$U_n^u = O(n^{1+\frac{1}{\beta}})$	$U_n^g = O(n^{1+\frac{1}{2\beta}})$	$z_n^* = \frac{\lambda}{\eta(1+\beta)} n$ $U_n^t(z_n^*) = O(n^{1+\frac{1}{\beta}})$ does not converge to upper bound

Simulation Results

Utility Distribution	Optimal Threshold		Simulated Utility Rate		
	Theory	Simulation	Theoretical Threshold	Simulation Threshold	Greedy Policy
exp(1)	144.8	148	4833	4833	3462

Scenario: $\lambda = \eta = 1$, $n = 1000$

Simulation Results

Utility Distribution	Optimal Threshold		Simulated Utility Rate		
	Theory	Simulation	Theoretical Threshold	Simulation Threshold	Greedy Policy
exp(1)	144.8	148	4833	4833	3462
Pareto(1,2)	333.3	347	22,095	22,102	8259

Scenario: $\lambda = \eta = 1$, $n = 1000$

Simulation Results

Utility Distribution	Optimal Threshold		Simulated Utility Rate		
	Theory	Simulation	Theoretical Threshold	Simulation Threshold	Greedy Policy
exp(1)	144.8	148	4833	4833	3462
Pareto(1,2)	333.3	347	22,095	22,102	8259
U[0,1]	0	22	908.4	946.3	908.4

Scenario: $\lambda = \eta = 1$, $n = 1000$

Fluid Model for Population-Based Threshold Policy

$$
\begin{aligned}
B_n(t) &= B_n(0) + \int_0^t I_{\{S_n(r_-) < z_n\}} dN_B^+(\lambda n r) - N_B^-\left(\eta \int_0^t B_n(r)\,dr\right) \\
&\quad - \int_0^t I_{\{B_n(r_-) \geqslant z_n\}} dN_S^+(\lambda n r), \\
S_n(t) &= S_n(0) + \int_0^t I_{\{B_n(r_-) < z_n\}} dN_S^+(\lambda n r) - N_S^-\left(\eta \int_0^t S_n(r)\,dr\right) \\
&\quad - \int_0^t I_{\{S_n(r_-) \geqslant z_n\}} dN_B^+(\lambda n r).
\end{aligned}
$$

Fluid Model for Population-Based Threshold Policy

$$
\begin{aligned}
B_n(t) = {} & B_n(0) + \int_0^t I_{\{S_n(r_-)<z_n\}} dN_B^+(\lambda n r) - N_B^-\left(\eta \int_0^t B_n(r)\, dr\right) \\
& - \int_0^t I_{\{B_n(r_-)\geqslant z_n\}} dN_S^+(\lambda n r), \\
S_n(t) = {} & S_n(0) + \int_0^t I_{\{B_n(r_-)<z_n\}} dN_S^+(\lambda n r) - N_S^-\left(\eta \int_0^t S_n(r)\, dr\right) \\
& - \int_0^t I_{\{S_n(r_-)\geqslant z_n\}} dN_B^+(\lambda n r).
\end{aligned}
$$

Let $\bar{B}_n(t) = \frac{B_n(t)}{n}$ and $\bar{S}_n(t) = \frac{S_n(t)}{n}$ and let $n \to \infty$

Fluid Model for Population-Based Threshold Policy

$$B_n(t) = B_n(0) + \int_0^t I_{\{S_n(r_-)<z_n\}} dN_B^+(\lambda n r) - N_B^-\left(\eta \int_0^t B_n(r)\,dr\right)$$
$$- \int_0^t I_{\{B_n(r_-)\geqslant z_n\}} dN_S^+(\lambda n r),$$
$$S_n(t) = S_n(0) + \int_0^t I_{\{B_n(r_-)<z_n\}} dN_S^+(\lambda n r) - N_S^-\left(\eta \int_0^t S_n(r)\,dr\right)$$
$$- \int_0^t I_{\{S_n(r_-)\geqslant z_n\}} dN_B^+(\lambda n r).$$

Let $\bar{B}_n(t) = \frac{B_n(t)}{n}$ and $\bar{S}_n(t) = \frac{S_n(t)}{n}$ and let $n \to \infty$

$$\bar{B}(t) = \bar{B}(0) + \lambda t - \eta \int_0^t \bar{B}(r)\,dr - \lambda \int_0^t I_{\{\bar{B}(r)\geqslant z\}} dr - \lambda \int_0^t I_{\{\bar{S}(r)\geqslant z\}} dr$$
$$\bar{S}(t) = \bar{S}(0) + \lambda t - \eta \int_0^t \bar{S}(r)\,dr - \lambda \int_0^t I_{\{\bar{B}(r)\geqslant z\}} dr - \lambda \int_0^t I_{\{\bar{S}(r)\geqslant z\}} dr$$

Utility-based Threshold Policy With Threshold v_n

- If $\max\{V_1, \ldots, V_{S_n(t)}\} > v_n$ then buyer matches upon arrival
- If $\max\{V_1, \ldots, V_{S_n(t)}\} \leqslant v_n$ then buyer waits in market

Utility-based Threshold Policy With Threshold v_n

- If $\max\{V_1, \ldots, V_{S_n(t)}\} > v_n$ then buyer matches upon arrival
- If $\max\{V_1, \ldots, V_{S_n(t)}\} \leqslant v_n$ then buyer waits in market
- Similar match/wait decision when seller arrives to find $B_n(t)$ buyers

Utility-based Threshold Policy With Threshold v_n

- $\left(\frac{B_n(t)}{n}, \frac{S_n(t)}{n}\right)$ converges to fluid limit (Ethier and Kurtz 1986)
- Pair of ODEs describing fluid limit has a unique stationary point

Utility-based Threshold Policy With Threshold v_n

- $\left(\frac{B_n(t)}{n}, \frac{S_n(t)}{n}\right)$ converges to fluid limit (Ethier and Kurtz 1986)
- Pair of ODEs describing fluid limit has a unique stationary point
- *Theorem*: Under some technical assumptions, a threshold of the form $v_n = v_* E[M_n]$ with $v_* > 0$ suitably chosen is asymptotically optimal within the class of utility-based threshold policies

Utility-based Threshold Policy With Threshold v_n

- $\left(\frac{B_n(t)}{n}, \frac{S_n(t)}{n}\right)$ converges to fluid limit (Ethier and Kurtz 1986)
- Pair of ODEs describing fluid limit has a unique stationary point
- *Theorem*: Under some technical assumptions, a threshold of the form $v_n = v_* E[M_n]$ with $v_* > 0$ suitably chosen is asymptotically optimal within the class of utility-based threshold policies
- For Pareto distribution ($\beta > 1$), optimal threshold reduces to optimizing an expression involving the Lambert W function
- For Pareto(1,2) example (mean utility $= 2$), the optimal computed threshold is 42.8

Population-based Threshold vs. Utility-based Threshold

Utility Dist'n	Population-based Threshold			Utility-based Threshold		
	Optimal Threshold	Utility Rate	Fraction Abandon	Optimal Threshold	Utility Rate	Fraction Abandon
exp(1)	148	4833	0.140	5.6	5732	0.150

Scenario: $\lambda = \eta = 1$, $n = 1000$

Population-based Threshold vs. Utility-based Threshold

Utility Dist'n	Population-based Threshold			Utility-based Threshold		
	Optimal Threshold	Utility Rate	Fraction Abandon	Optimal Threshold	Utility Rate	Fraction Abandon
exp(1)	148	4833	0.140	5.6	5732	0.150
Pareto (1,2)	347	22,102	0.334	42.0	43,750	0.503

Scenario: $\lambda = \eta = 1,\ n = 1000$

Population-based Threshold vs. Utility-based Threshold

Utility Dist'n	Population-based Threshold			Utility-based Threshold		
	Optimal Threshold	Utility Rate	Fraction Abandon	Optimal Threshold	Utility Rate	Fraction Abandon
exp(1)	148	4833	0.140	5.6	5732	0.150
Pareto (1,2)	347	22,102	0.334	42.0	43,750	0.503
U[0,1]	22	946	0.027	0.96	963	0.021

Scenario: $\lambda = \eta = 1$, $n = 1000$

Fluid Model for Utility-Based Threshold Policy

$$B_n(t) = B_n(0) + \sum_{j=1}^{N_B^+(n\lambda t)} I_{\left\{\max_{i=1}^{S_n\left(A_{j-}^B\right)} V_{i,j}^B \leqslant v\right\}} - \sum_{j=1}^{N_S^+(n\lambda t)} I_{\left\{\max_{i=1}^{B_n(A_{j-}^S)} V_{i,j}^S > v\right\}}$$

$$- N_B^-\left(\eta \int_0^t B_n(r_-)\, dr\right)$$

Fluid Model for Utility-Based Threshold Policy

$$B_n(t) = B_n(0) + \sum_{j=1}^{N_B^+(n\lambda t)} I_{\left\{\max_{i=1}^{S_n\left(A_{j-}^B\right)} V_{i,j}^B \leqslant v\right\}} - \sum_{j=1}^{N_S^+(n\lambda t)} I_{\left\{\max_{i=1}^{B_n\left(A_{j-}^S\right)} V_{i,j}^S > v\right\}} - N_B^-\left(\eta \int_0^t B_n(r_-)\,dr\right)$$

Let $\bar{B}_n(t) = \frac{B_n(t)}{n}$

$$\bar{B}_n(t) = \bar{B}_n(0) + \frac{N_B^+\left(\lambda n \int_0^t F_{S_n(r)}(v_n)\,dr\right)}{n} - \frac{N_B^-\left(\eta \int_0^t B_n(r)\,dr\right)}{n} - \frac{\tilde{N}_S^+\left(\lambda n \int_0^t \left(1 - F_{B_n(r)}(v_n)\right)dr\right)}{n}$$

$$\bar{B}_n(t) = \bar{B}_n(0) + \frac{N_B^+\left(\lambda n \int_0^t F_{S_n(r)}(v_n)\,dr\right)}{n} - \frac{N_B^-\left(\eta \int_0^t B_n(r)\,dr\right)}{n} - \frac{\tilde{N}_S^+\left(\lambda n \int_0^t \left(1 - F_{B_n(r)}(v_n)\right) dr\right)}{n}$$

Let $n \to \infty$. Under Assumption 3,

$$\bar{B}(t) = \bar{B}(0) + \lambda \int_0^t e^{-\kappa \bar{S}(r)/v^{1/\alpha}} - \eta \int_0^t \bar{B}(r)\,dr - \lambda \int_0^t \left(1 - e^{-\kappa \bar{B}(r)/v^{1/\alpha}}\right) dr,$$

Fluid Model for Utility-Based Threshold Policy

$$\bar{B}_n(t) = \bar{B}_n(0) + \frac{N_B^+\left(\lambda n \int_0^t F_{S_n(r)}(v_n)\, dr\right)}{n} - \frac{N_B^-\left(\eta \int_0^t B_n(r)\, dr\right)}{n} - \frac{\tilde{N}_S^+\left(\lambda n \int_0^t \left(1 - F_{B_n(r)}(v_n)\right) dr\right)}{n}$$

Let $n \to \infty$. Under Assumption 3,

$$\bar{B}(t) = \bar{B}(0) + \lambda \int_0^t e^{-\kappa \bar{S}(r)/v^{1/\alpha}} - \eta \int_0^t \bar{B}(r)\, dr - \lambda \int_0^t \left(1 - e^{-\kappa \bar{B}(r)/v^{1/\alpha}}\right) dr,$$

Solve $0 = \lambda e^{-\kappa \bar{z}/v^{1/\alpha}} - \eta \bar{z} - \lambda(1 - e^{-\kappa \bar{z}/v^{1/\alpha}})$

Correlated Utilities

- Under population-threshold policy, the optimal threshold is independent of correlation ρ and the utility rate is decreasing in ρ
- Under utility-threshold policy, the optimal threshold is decreasing in ρ and the utility rate is decreasing in ρ

Unbalanced Markets

- $\lambda_B \neq \lambda_S$ and $\eta_B \neq \eta_S$ with heavy tails
- Buyers and sellers have the same utility-based threshold
- Market thickness (i.e., threshold value) increases with amount of imbalance

Batch and Match

- Can we do better if we drop the requirement that matches must be made at arrival epochs (and involve an arriving item)?

Batch and Match

- Consider a policy that:
  - Collects all arrivals over a time interval of length Δ
  - Chooses at random $\min\{B(t), S(t)\}$ agents from thicker side of market
  - Maximizes utility from these matches

Batch and Match

- Consider a policy that:
  - Collects all arrivals over a time interval of length Δ
  - Chooses at random $\min\{B(t), S(t)\}$ agents from thicker side of market
  - Maximizes utility from these matches

- *Theorem:* Under a Pareto(c, β) distribution with finite mean, the utility rate $U_n^b(\Delta)$ satisfies

$$\lim_{n \to \infty} \frac{U_n^b(\Delta)}{n^{\alpha+1}} \leqslant f(c, \alpha, \lambda, \eta, \Delta^*)$$

where the optimal time window Δ^* is the unique solution to

$$e^{\eta\Delta} = (1+\alpha)\eta\Delta + 1$$

Batch and Match

- Consider a policy that:
  - Collects all arrivals over a time interval of length Δ
  - Chooses at random $\min\{B(t), S(t)\}$ agents from thicker side of market
  - Maximizes utility from these matches
- Let $\lambda = \eta = 1$, $n = 1000$, $c = 1, \beta = 2$:
  - Upper bound for batch utility is less than utility from utility threshold policy
  - In simulations, batch utility $= 25$k (vs 44k for utility threshold policy)
  - Average batch size $= 532$ matches

Summary

- We incorporated general utilities into centralized dynamic matching markets by using
  - Queueing asymptotics
  - Extreme value theory

Summary

- We incorporated general utilities into centralized dynamic matching markets by using
  - Queueing asymptotics
  - Extreme value theory
- As right tail of matching utility distribution gets heavier:
  - Optimal market thickness increases
  - Abandonment increases
  - Optimal utility rate increases

Summary

- We incorporated general utilities into centralized dynamic matching markets by using
  - Queueing asymptotics
  - Extreme value theory
- As right tail of matching utility distribution gets heavier:
  - Optimal market thickness increases
  - Abandonment increases
  - Optimal utility rate increases
- Empirical work (Hitsch et al. 2010, Boyd et al. 2013, Agrawal 2015) suggests that large centralized matching markets are likely to benefit from allowing the market to thicken