Matching in Dynamic Imbalanced Markets

Itai Ashlagi Afshin Nikzad Philipp Strack

Intro

Research Question:
How fast should one match (clear the market)?

Intro

Research Question:
How fast should one match (clear the market)?

Main Trade-Off:

1. Matching faster reduces waiting time (often costly, i.e. time on dialysis)
2. Waiting “thickens” the market and might facilitate more and better matches.

Intro

Research Question:
How fast should one match (clear the market)?

Main Trade-Off:

1. Matching faster reduces waiting time (often costly, i.e. time on dialysis)
2. Waiting "thickens" the market and might facilitate more and better matches.

$\rightarrow$ If agents are **heterogeneous** some might benefit from waiting while other don't.

Intro

Research Question:
How fast should one match (clear the market)?

Main Trade-Off:

1. Matching faster reduces waiting time (often costly, i.e. time on dialysis)
2. Waiting "thickens" the market and might facilitate more and better matches.

→ If agents are **heterogeneous** some might benefit from waiting while other don't.
  - → In kidney exchanges matching fast might hurt hard-to-match agents.

Intro

Research Question:
How fast should one match (clear the market)?

Main Trade-Off:

1. Matching faster reduces waiting time (often costly, i.e. time on dialysis)
2. Waiting "thickens" the market and might facilitate more and better matches.

→ If agents are **heterogeneous** some might benefit from waiting while other don't.
  → In kidney exchanges matching fast might hurt hard-to-match agents.

→ Potentially relevant in many context: ride-sharing, kidneys, financial markets, ..

Intro

Research Question:
How fast should one match (clear the market)?

Main Trade-Off:

1. Matching faster reduces waiting time (often costly, i.e. time on dialysis)
2. Waiting "thickens" the market and might facilitate more and better matches.

$\rightarrow$ If agents are **heterogeneous** some might benefit from waiting while other don't.

  $\rightarrow$ In kidney exchanges matching fast might hurt hard-to-match agents.

$\rightarrow$ Potentially relevant in many context: ride-sharing, kidneys, financial markets, ..

$\rightarrow$ We look at this question in **large kidney exchanges** without match quality.

Kidney Exchanges

- About 18000 transplant per year:

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;

Kidney Exchanges

- About 18000 transplant per year:
    - 12000 from cadaver organs;
    - 6000 from living donors;

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant (“too sick”) while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.
  - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).

Kidney Exchanges

- About 18000 transplant per year:
    - 12000 from cadaver organs;
    - 6000 from living donors;
    - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:
    - Direct donations: patient and compatible donor.
    - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).
- Exchanges match at different speeds:

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.
  - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).
- Exchanges match at different speeds:
  - US: gradually more frequenetly. Daily (NKR, APD, Methodist in SA), bi-weekly (UNOS)

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.
  - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).
- Exchanges match at different speeds:
  - US: gradually more frequenetly. Daily (NKR, APD, Methodist in SA), bi-weekly (UNOS)
  - Europe (Netherlands, UK, Czech Republic), Canada, Australia: 3-4 months.

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant ("too sick") while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.
  - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).
- Exchanges match at different speeds:
  - US: gradually more frequenetly. Daily (NKR, APD, Methodist in SA), bi-weekly (UNOS)
  - Europe (Netherlands, UK, Czech Republic), Canada, Australia: 3-4 months.
  - Israel: daily.

Kidney Exchanges

- About 18000 transplant per year:
  - 12000 from cadaver organs;
  - 6000 from living donors;
  - But 5000 are removed without a transplant (“too sick”) while 35000 joined the list.
- Live donation helps to overcome shortage:
  - Direct donations: patient and compatible donor.
  - Exchange: between incompatible patient-donor pairs ($\approx$ 15% of US live donor transplants).
- Exchanges match at different speeds:
  - US: gradually more frequenetly. Daily (NKR, APD, Methodist in SA), bi-weekly (UNOS)
  - Europe (Netherlands, UK, Czech Republic), Canada, Australia: 3-4 months.
  - Israel: daily.
- Concerns that high matching frequency in the US leads to inefficiency: “There has been a race to the bottom in that registries forced by competition to perform match runs very frequently...and likely fewer transplants are accomplished nationwide” (Gentry and Segev, AJT 2015).

Related literature

- Minimize waiting times (with no exogenous departures): Unver 2010, Gurvich and Ward 2016, Anderson, Ashlagi, Gamarnik Kanoria 2017 , Ashlagi, Burq, Manshadi, Jaillet 2019, Blum and Mansour 2020, Akbarpour, Combe, He, Hiller, Shimer, Tercieux 2020.
- Maximize number of matches with departures: Akbarpour, Li, Oveis-Gharan 2020, Nikzad, Akbarpour, Rees, Roth 2020
- Simulations: Ashlagi et al. (AJT) 2017, Agarwal, Ashlagi, Azevedo, Featherstone, Karaduman 2018.
- Dynamic matching with heterogeneous values: Doval 2014. Liu, Wang, Yang, 2018, Baccara, Lee, Yariv 2019, Mertikopoulos, Nax, Pradelski 2019, Blanchet, Reiman, Shah, Wein 2020, Collina, Immorlica, Leton-Brown, Lucier, 2020.
- Batching and clearing times: Mendhelson 1982, Loertscher, Taylor, Muir, 2016, Li et al. 2019, Matteo, Budish, Oneil 2020, Kerimov, Ashlagi, Gurvich 2021.
- Dynamic programming for KE: Dickerson, Procaccia, Sandholm 2012a,b

Overview

Overview

- **Analysis of the compatibility graph.**

Overview

- **Analysis of the compatibility graph.**
    1. data: Identify empirical features of large compatibility graphs .

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.
- **Analysis of dynamic matching policies.**

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.
- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.
- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.
- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.
- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.
     - Patient is suboptimal in large markets.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.
     - Patient is suboptimal in large markets.
  3. Empirical analysis:

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.
     - Patient is suboptimal in large markets.
  3. Empirical analysis:
     - Greedy does better than 7d, 14d, 30d batching for moderate market sizes both in simulations of the model and with real compatibility data.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.
     - Patient is suboptimal in large markets.
  3. Empirical analysis:
     - Greedy does better than 7d, 14d, 30d batching for moderate market sizes both in simulations of the model and with real compatibility data.
     - Patient matches around 1% more, but leads to 35% longer waiting time.

Overview

- **Analysis of the compatibility graph.**
  1. data: Identify empirical features of large compatibility graphs .
  2. theory: Single type random graph models can not fit these features.
  3. theory: Propose a simple two-type model.

- **Analysis of dynamic matching policies.**
  1. Consider three policies: (1) Greedy matching (2) Batching (3) Patient matching.
  2. Theoretical analysis:
     - Greedy matching is (almost) optimal in large markets *for all types* and *all linear EU preferences*.
     - Batching policies can only be optimal they match at high frequency.
     - Patient is suboptimal in large markets.
  3. Empirical analysis:
     - Greedy does better than 7d, 14d, 30d batching for moderate market sizes both in simulations of the model and with real compatibility data.
     - Patient matches around 1% more, but leads to 35% longer waiting time.
     - Interestingly, patient matching is bad for hard-to-match agents.

The compatibility graph

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.
- Compatibility graph G: agent $\equiv$ node, edge $\equiv$ compatibility.

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.
- Compatibility graph G: agent $\equiv$ node, edge $\equiv$ compatibility.
- Matching $\mu \in M(G)$: pairwise assignment of compatible agents.

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.
- Compatibility graph G: agent $\equiv$ node, edge $\equiv$ compatibility.
- Matching $\mu \in M(G)$: pairwise assignment of compatible agents.
- **Size of the maximum matching**:

$$\mathrm{SMM} = \max_{\mu \in M(G)} \frac{|\mu|}{|G|} .$$

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.
- Compatibility graph G: agent $\equiv$ node, edge $\equiv$ compatibility.
- Matching $\mu \in M(G)$: pairwise assignment of compatible agents.
- **Size of the maximum matching**:

$$\text{SMM} = \max_{\mu \in M(G)} \frac{|\mu|}{|G|}.$$

- **Fraction of agents without a partner**:

$$\text{FWP} = \frac{|\{i \in G : (i,j) \notin M(G) \text{ for all } j\}|}{|G|}.$$

Compatibility Graphs

- An agent is a pair of a kidney donor and a patient.
- Two agents are compatibly if they can exchange kidneys.
- Compatibility graph G: agent $\equiv$ node, edge $\equiv$ compatibility.
- Matching $\mu \in M(G)$: pairwise assignment of compatible agents.
- **Size of the maximum matching**:

$$\text{SMM} = \max_{\mu \in M(G)} \frac{|\mu|}{|G|}.$$

- **Fraction of agents without a partner**:

$$\text{FWP} = \frac{|\{i \in G : (i,j) \notin M(G) \text{ for all } j\}|}{|G|}.$$

- Capture the benefit of enlarging the market exogenously or by waiting.

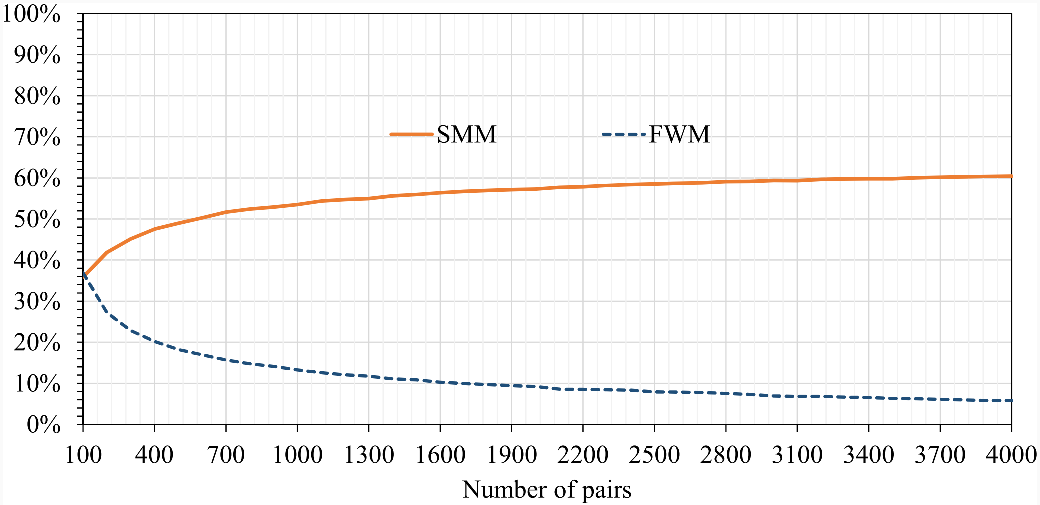

Figure 1: Average percentage of pairs without a compatible partner (dashed) and the percentage matched in a maximum matching (solid). The average for every fixed pool size on the horizontal axis is computed by random sampling from the combined data set from NKR, APD, UNOS and Methodist at San Antonio.

Empirical Regularities

Two empirical regularities:

1. Size of the maximum matching is bounded away from 1.
2. Fraction of agents without a partner goes to 0.

Empirical Regularities

Two empirical regularities:

1. Size of the maximum matching is bounded away from 1.
2. Fraction of agents without a partner goes to 0.

Proposition

Consider a model with m homogeneous agents, in which every pair of agents are compatible independently with probability $p(m) > 0$ that may depend on the market size. The following two conditions cannot be satisfied simultaneously

$$\lim_{m \to \infty} \mathbb{E}\left[\mathrm{SMM}\right] < 1, \textit{ and} \tag{1}$$

$$\lim_{m \to \infty} \mathbb{E}\left[\mathrm{FWP}\right] = 0\,. \tag{2}$$

Empirical Regularities

Two empirical regularities:

1. Size of the maximum matching is bounded away from 1.
2. Fraction of agents without a partner goes to 0.

Proposition

Consider a model with m homogeneous agents, in which every pair of agents are compatible independently with probability $p(m) > 0$ that may depend on the market size. The following two conditions cannot be satisfied simultaneously

$$\lim_{m\to\infty} \mathbb{E}\left[\mathrm{SMM}\right] < 1, \textit{ and} \tag{1}$$

$$\lim_{m\to\infty} \mathbb{E}\left[\mathrm{FWP}\right] = 0\,. \tag{2}$$

→ Intuitively, heterogeneity plays a major role.

A simple two-type model

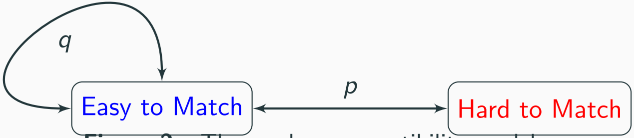

Figure 2: The random compatibility model.

A simple two-type model

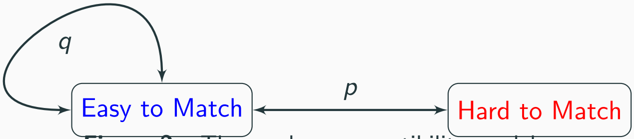

Figure 2: The random compatibility model.

- Two types easy (E) and hard-to-match (H).
- Compatibility is independent across pairs
  1. $p > 0$ between (E) and (H);
  2. $q > 0$ between (E) and (E);
  3. 0 between (H) and (H);

Relating this model to data

Proposition

Consider the compatibility model with m easy-to-match agents and $(1+\lambda)m$ hard-to-match agents where $\lambda > 0$. Then, with high probability[1] *we have that*

$$\mathrm{SMM} = \frac{2}{2+\lambda}, \tag{3}$$

$$\mathrm{FWP} = 0\,. \tag{4}$$

[1] A sequence of events $E_1, E_2, \ldots$ holds with high probability if there is $\alpha > 0$ such that $\lim_{n\to\infty} n^{\alpha}(1 - \mathbb{P}[E_n]) = 0$.

Relating this model to data

Proposition

Consider the compatibility model with m easy-to-match agents and $(1+\lambda)m$ hard-to-match agents where $\lambda > 0$. Then, with high probability[1] *we have that*

$$\mathrm{SMM} = \frac{2}{2+\lambda}, \tag{3}$$

$$\mathrm{FWP} = 0\,. \tag{4}$$

- There are λm more H agents which go unmatched $\Rightarrow \mathrm{SMM} \leq \frac{2}{2+\lambda}$.

[1] A sequence of events $E_1, E_2, \ldots$ holds with high probability if there is $\alpha > 0$ such that $\lim_{n\to\infty} n^{\alpha}(1 - \mathbb{P}[E_n]) = 0$.

Relating this model to data

Proposition

Consider the compatibility model with m easy-to-match agents and $(1+\lambda)m$ hard-to-match agents where $\lambda > 0$. Then, with high probability[1] *we have that*

$$\mathrm{SMM} = \frac{2}{2+\lambda}, \tag{3}$$

$$\mathrm{FWP} = 0\,. \tag{4}$$

- There are λm more H agents which go unmatched $\Rightarrow \mathrm{SMM} \leq \frac{2}{2+\lambda}$.
- Perfect matching in bipartite with high prob. $\Rightarrow \mathrm{SMM} = \frac{2}{2+\lambda}$ with high prob.

[1] A sequence of events $E_1, E_2, \ldots$ holds with high probability if there is $\alpha > 0$ such that $\lim_{n\to\infty} n^{\alpha}(1 - \mathbb{P}[E_n]) = 0$.

Relating this model to data

Proposition

Consider the compatibility model with m easy-to-match agents and $(1+\lambda)m$ hard-to-match agents where $\lambda > 0$. Then, with high probability[1] *we have that*

$$\mathrm{SMM} = \frac{2}{2+\lambda}, \tag{3}$$

$$\mathrm{FWP} = 0\,. \tag{4}$$

- There are λm more H agents which go unmatched $\Rightarrow \mathrm{SMM} \leq \frac{2}{2+\lambda}$.
- Perfect matching in bipartite with high prob. $\Rightarrow \mathrm{SMM} = \frac{2}{2+\lambda}$ with high prob.
- In the data $\lim_{m\to\infty} \mathrm{SMM} = 0.6$ suggesting $\lambda \approx 1.3$, i.e. 70% hard-to-match.

[1] A sequence of events $E_1, E_2, \dots$ holds with high probability if there is $\alpha > 0$ such that $\lim_{n\to\infty} n^{\alpha}(1 - \mathbb{P}[E_n]) = 0$.

Dynamic Matching

Dynamic Matching

- Agents **arrive** according to Poisson process.

Dynamic Matching

- Agents **arrive** according to Poisson process.
    1. H agents arrive at rate $m(1+\lambda)$.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.

Dynamic Matching

- Agents **arrive** according to Poisson process.
    1. H agents arrive at rate $m(1+\lambda)$.
    2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
    1. Criticality is observable.
    2. Last time an agent can match.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.
  2. Last time an agent can match.

- A dynamic policy chooses at each point in time a matching $\mu_t \in M(G_t)$ to execute.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.
- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.
  2. Last time an agent can match.
- A dynamic policy chooses at each point in time a matching $\mu_t \in M(G_t)$ to execute.
  1. **Greedy**: execute every possible matching immediately.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.
  2. Last time an agent can match.

- A dynamic policy chooses at each point in time a matching $\mu_t \in M(G_t)$ to execute.
  1. **Greedy**: execute every possible matching immediately.
  2. **Batching**: every T units of time execute a maximal matching.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.
  2. Last time an agent can match.

- A dynamic policy chooses at each point in time a matching $\mu_t \in M(G_t)$ to execute.
  1. **Greedy**: execute every possible matching immediately.
  2. **Batching**: every T units of time execute a maximal matching.
  3. **Patient**: whenever an agent gets critical attempt to match that agent.

Dynamic Matching

- Agents **arrive** according to Poisson process.
  1. H agents arrive at rate $m(1+\lambda)$.
  2. E agents arrive at rate m.

- Agents **become critical** at an exponentially distributed time with mean d.
  1. Criticality is observable.
  2. Last time an agent can match.

- A dynamic policy chooses at each point in time a matching $\mu_t \in M(G_t)$ to execute.
  1. **Greedy**: execute every possible matching immediately.
  2. **Batching**: every T units of time execute a maximal matching.
  3. **Patient**: whenever an agent gets critical attempt to match that agent.

  Throughout break ties randomly in favor of H agents.

Why can waiting be beneficial?

Time 1:

Why can waiting be beneficial?

Time 1:

Time 2:

Why can waiting be beneficial?

Time 1:

Time 2:

- The policy which matches at time 1 matches 2 agents.

Why can waiting be beneficial?

Time 1:

Time 2:

- The policy which matches at time 1 matches 2 agents.
- The policy which matches at time 2 matches 4 agents.

Measure for performance

- $\theta_i \in \{E, H\}$ agent i's type
- $\alpha_i \geq 0$ her arrival time
- $\varphi_i \geq 0$ how long she is present in the market
- $\mu_i \in \{0, 1\}$ whether she is matched.

Measure for performance

- $\theta_i \in \{E, H\}$ agent i's type
- $\alpha_i \geq 0$ her arrival time
- $\varphi_i \geq 0$ how long she is present in the market
- $\mu_i \in \{0, 1\}$ whether she is matched.

- **Match rate**

$$q_\Theta(m) = \lim_{t \to \infty} \mathbb{E}\left[\frac{|\{i\colon \mu_i = 1 \text{ and } \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right].$$

- **Waiting time**

$$w_\Theta(m) = \lim_{t \to \infty} \mathbb{E}\left[\frac{\sum_{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta} \varphi_i}{|\{i\colon \alpha_i \leq t \text{ and } \theta_i = \Theta\}|}\right].$$

- Motivation: payoff of a risk-neutral expected-utility-maximizer with constant waiting cost.

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

- Demanding as the policy needs to be optimal

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

- Demanding as the policy needs to be optimal
  1. for *all* types.

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

- Demanding as the policy needs to be optimal
  1. for *all* types.
  2. for *all* risk-neutral preferences with linear waiting cost.

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

- Demanding as the policy needs to be optimal
  1. for *all* types.
  2. for *all* risk-neutral preferences with linear waiting cost.
- Existence of such a policy is unclear.

Asymptotic Optimality

Definition (Asymptotic optimality)
A policy is asymptotically optimal if for every $\epsilon > 0$ there exists m_ϵ such that, when $m \geq m_\epsilon$, *no* type of agent can improve its match rate $q_\Theta(m)$ or expected waiting time $w_\Theta(m)$ by more than ϵ when changing to any other policy.

- Demanding as the policy needs to be optimal
  1. for *all* types.
  2. for *all* risk-neutral preferences with linear waiting cost.
- Existence of such a policy is unclear.
- If such a policy exists it there is no conflict between different types in a large market.

Results

Theorem

The greedy policy is asymptotically optimal, whereas the batching policy (for any fixed batch length) and the patient policy are not asymptotically optimal.

Results

Theorem

The greedy policy is asymptotically optimal, whereas the batching policy (for any fixed batch length) and the patient policy are not asymptotically optimal.

- Greedy is (almost) optimal for H **and** E agents in sufficiently large markets.
- Patient and Batching with fixed batch length are suboptimal in sufficiently large markets.

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:
  1. For $T \to 0$ Batching converges to Greedy (for a fixed market size).

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:
  1. For $T \to 0$ Batching converges to Greedy (for a fixed market size).
  2. For $T > 0$ matches both fewer (H, E) agents and matches the slower.

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:
  1. For $T \to 0$ Batching converges to Greedy (for a fixed market size).
  2. For $T > 0$ matches both fewer (H, E) agents and matches the slower.
- Patient:

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:
  1. For $T \to 0$ Batching converges to Greedy (for a fixed market size).
  2. For $T > 0$ matches both fewer (H, E) agents and matches the slower.
- Patient:
  1. Achieves the optimal match rates.

Proposition: As the arrival rate m grows large match rate and waiting time converge to:

	Matchrate q		Waiting time w	
	H	E	H	E
Greedy	$\frac{1}{1+\lambda}$	1	$\frac{\lambda d}{1+\lambda}$	0
Batching	$\frac{1-e^{-T/d}}{(1+\lambda)T/d}$	$\frac{1-e^{-T/d}}{T/d}$	$d(1-q_H)$	$d(1-q_E)$
Patient	$\frac{1}{1+\lambda}$	1	d	0

- d expected criticality time
- λ imbalance parameter
- T batching time

- Batching:
  1. For $T \to 0$ Batching converges to Greedy (for a fixed market size).
  2. For $T > 0$ matches both fewer (H, E) agents and matches the slower.

- Patient:
  1. Achieves the optimal match rates.
  2. Induces longer waiting times for H agents, but not for E.

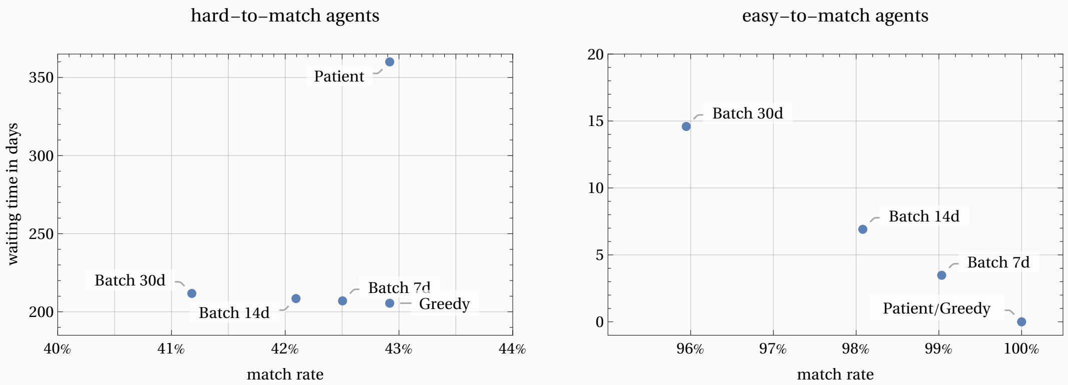

Figure 3: Illustration when $\lambda = 1.33$ and d equals 360 days. The blue points represent the predictions of our model for large markets which we derived.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.
- As H agents accumulate E agents can always always find a partner immediately.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.
- As H agents accumulate E agents can always always find a partner immediately.
- Under Greedy this achieves the upper bound in a large market.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.
- As H agents accumulate E agents can always always find a partner immediately.
- Under Greedy this achieves the upper bound in a large market.
- Under Batching the match rate is lower as some agents leave between matching intervals.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.
- As H agents accumulate E agents can always always find a partner immediately.
- Under Greedy this achieves the upper bound in a large market.
- Under Batching the match rate is lower as some agents leave between matching intervals.
- Under Patient H agents get only matched when they get critical and thus wait a long time.

Proof Intuition

- Match rates of $(\frac{1}{1+\lambda}, 1)$ are achieved by matching all E agents to H agents and thus constitute an upper bound.
- Clearly, a waiting time of 0 is a bound for the E waiting time.
- The upper bound on the waiting time for H agent is obtained by analyzing the process where all H and E agents are compatible.
- Intuitively, H agents must accumulate in any mechanism as they can only match to E agents and there are fewer E agents.
- As H agents accumulate E agents can always always find a partner immediately.
- Under Greedy this achieves the upper bound in a large market.
- Under Batching the match rate is lower as some agents leave between matching intervals.
- Under Patient H agents get only matched when they get critical and thus wait a long time.
- **Proofs:** Detailed analysis of the 2-dimensional Markov chain for each process.

How large do is large?

Greedy and Batching in Finite Markets

Proposition: A market size dependent batching policy with batch length T_m is asymptotically optimal if and only if the batch length goes to zero as the market becomes large $\lim_{m\to\infty} T_m = 0$.

Greedy and Batching in Finite Markets

Proposition: A market size dependent batching policy with batch length T_m is asymptotically optimal if and only if the batch length goes to zero as the market becomes large $\lim_{m\to\infty} T_m = 0$.

For how large batching time is batching suboptimal at a **fixed market size**?

Greedy and Batching in Finite Markets

Proposition: A market size dependent batching policy with batch length T_m is asymptotically optimal if and only if the batch length goes to zero as the market becomes large $\lim_{m\to\infty} T_m = 0$.

For how large batching time is batching suboptimal at a **fixed market size**?

Proposition: Let $m > 0$ be an arbitrary *fixed* arrival rate. Define z^* to be the steady-state probability that an E agent, upon her arrival, is matched to an H agent under the greedy policy. Then, for every E and H agents the match rate and waiting time of that type under the batching policy are worse than under the greedy policy if

$$T > \frac{z^* W\left(-\frac{e^{-1/z^*}}{z^*}\right) + 1}{z^*/d} \tag{5}$$

and $W(\cdot)$ is the Lambert W function.

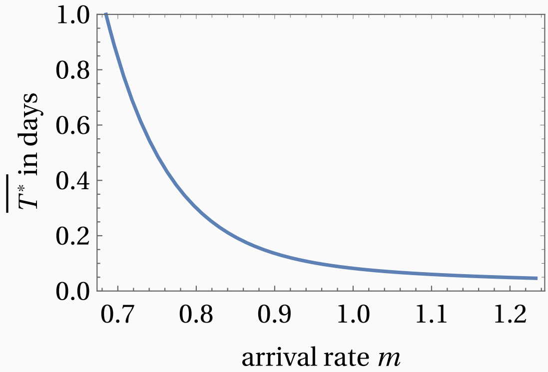

Figure 4: The batch length above which greedy dominates batching for various arrival rates per day, $\lambda = 1.33, p = 0.037$ and average criticality time $d = 360$ days. The bound $\overline{T^*}$ is independent of $q \in [0, 1]$, and is decreasing in p.

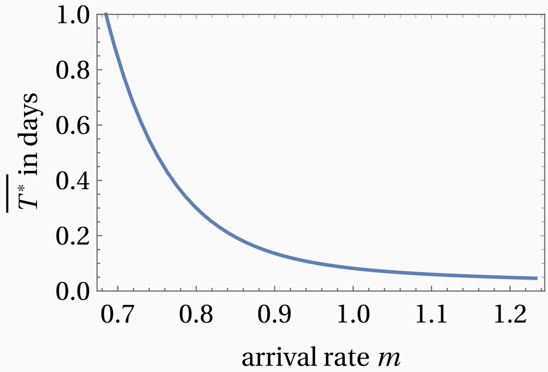

Figure 4: The batch length above which greedy dominates batching for various arrival rates per day, $\lambda = 1.33, p = 0.037$ and average criticality time $d = 360$ days. The bound $\overline{T^*}$ is independent of $q \in [0, 1]$, and is decreasing in p.

For 1.6 pairs arriving per day matching batching less frequently **than** daily is strictly sub-optimal for all types (National Kidney Registry $\equiv 1$ pair per day).

Model Simulations

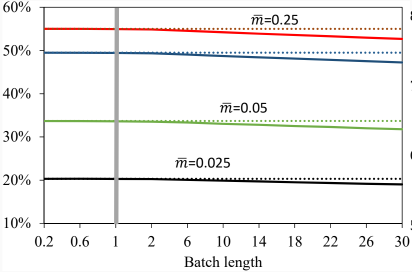

(a) Different arrival rates

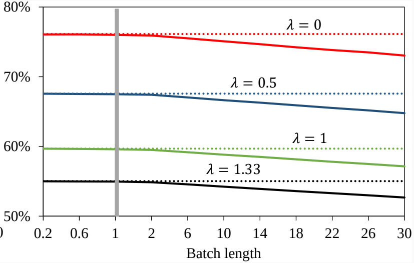

(b) Different imbalances

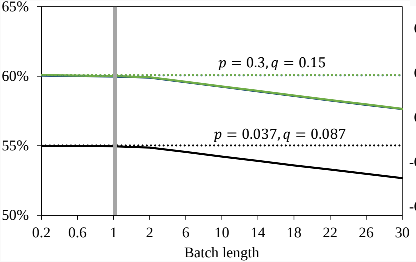

(c) Different compatibilities

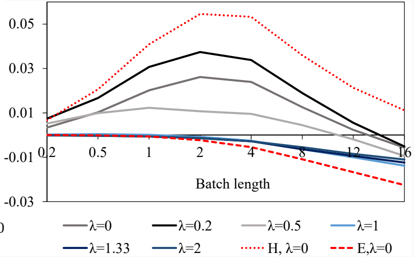

(d) Match probabilities $p = 0.02, q = 1$

Data Simulations

Simluations

- We use compatibility date from 1881 de-identified patient-donor pairs from the NKR (between July 2007 to December 2014).
- Arrivals and departures follow our model ($d = 360$).
- We vary market size between $1/10$ and 4 times the size of the NKR.
- We simulate the arrival of 10 million pairs for each policy.

arrivals per day	*match rate*					*waiting time in days*				
	Greedy	Patient	Batching			Greedy	Patient	Batching		
			7 days	30 days	60 days			7 days	30 days	60 days
0.01	10.7%	11.9%	10.4%	9.9%	9.3%	322	355	322	324	326
0.05	22.4%	23.4%	22.2%	21.2%	20.2%	279	324	280	283	288
0.25	34.3%	35.4%	33.8%	32.6%	31.2%	237	298	238	243	248
0.5	38.5%	39.5%	38.0%	36.8%	35.2%	222	290	223	228	233
1	42.0%	43%	41.6%	40.2%	38.6%	209	283	210	215	221
2	45%	45.8%	44.5%	43.1%	41.5%	198	278	200	205	211
4	47.2%	48%	46.8%	45.4%	43.6%	190	274	192	196	207

Table 1: Match rate and average waiting time over all pairs in simulations using NKR data.

arrivals per day	*match rate*					*waiting time in days*				
	Greedy	Patient	Batching			Greedy	Patient	Batching		
			7 days	30 days	60 days			7 days	30 days	60 days
0.01	10.7%	11.9%	10.4%	9.9%	9.3%	322	355	322	324	326
0.05	22.4%	23.4%	22.2%	21.2%	20.2%	279	324	280	283	288
0.25	34.3%	35.4%	33.8%	32.6%	31.2%	237	298	238	243	248
0.5	38.5%	39.5%	38.0%	36.8%	35.2%	222	290	223	228	233
1	42.0%	43%	41.6%	40.2%	38.6%	209	283	210	215	221
2	45%	45.8%	44.5%	43.1%	41.5%	198	278	200	205	211
4	47.2%	48%	46.8%	45.4%	43.6%	190	274	192	196	207

Table 1: Match rate and average waiting time over all pairs in simulations using NKR data.

- Patient leads to 35% longer waiting times, but $\approx 1\%$ higher match rate.

arrivals per day	*match rate*					*waiting time in days*				
	Greedy	Patient	Batching			Greedy	Patient	Batching		
			7 days	30 days	60 days			7 days	30 days	60 days
0.01	10.7%	11.9%	10.4%	9.9%	9.3%	322	355	322	324	326
0.05	22.4%	23.4%	22.2%	21.2%	20.2%	279	324	280	283	288
0.25	34.3%	35.4%	33.8%	32.6%	31.2%	237	298	238	243	248
0.5	38.5%	39.5%	38.0%	36.8%	35.2%	222	290	223	228	233
1	42.0%	43%	41.6%	40.2%	38.6%	209	283	210	215	221
2	45%	45.8%	44.5%	43.1%	41.5%	198	278	200	205	211
4	47.2%	48%	46.8%	45.4%	43.6%	190	274	192	196	207

Table 1: Match rate and average waiting time over all pairs in simulations using NKR data.

- Patient leads to 35% longer waiting times, but $\approx 1\%$ higher match rate.
- Greedy does better than $7d, 30d, 60d$ Batching in terms of waiting times and match rate.

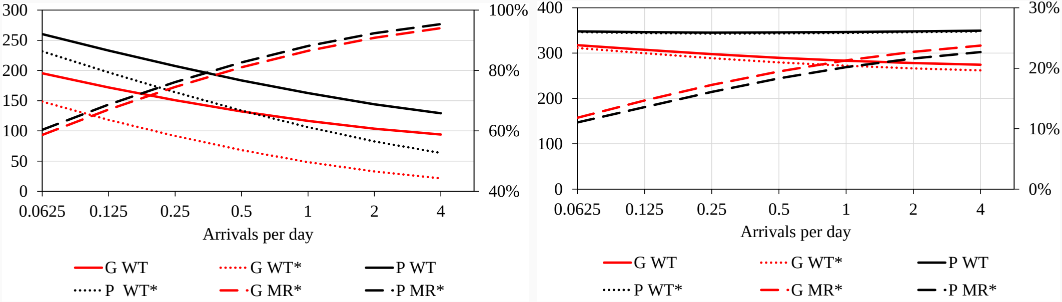

(a) Over-demanded pairs

(b) Under-demanded pairs

Figure 6: Average waiting times (WT) and match rate (MR) in days under greedy (G) and patient (P) policies. The left and right axes are WT and MR. The label (*) excludes pairs who have no match in the data. Under-demanded patient-donor pairs are blood type incompatible, Over-demanded pairs are blood type compatible, but tissue type incompatible.

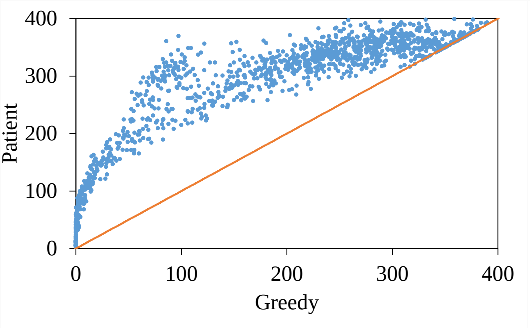

(a) Waiting times

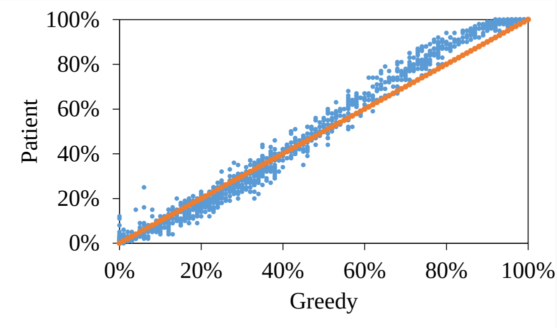

(b) Chance of matching

Figure 7: Averages of waiting times (left) and chance of matching (right) taken over copies for each pair in the data. The axes correspond to the greedy and patient policies. Arrival 1 per day.

Modelling Assumptions

Robustness

- Single Type.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.

Robustness

- Single Type.
    - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
    - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.

- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.
  - Simulations for $p = q = \frac{c}{m}$: Greedy has a lower waiting time and Patient a higher match rate.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.
  - Simulations for $p = q = \frac{c}{m}$: Greedy has a lower waiting time and Patient a higher match rate.
- Different Objectives.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.
  - Simulations for $p = q = \frac{c}{m}$: Greedy has a lower waiting time and Patient a higher match rate.
- Different Objectives.
  - We show that for *all* risk-neutral EU preferences with linear waiting cost Greedy is optimal in large markets.

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.
  - Simulations for $p = q = \frac{c}{m}$: Greedy has a lower waiting time and Patient a higher match rate.
- Different Objectives.
  - We show that for *all* risk-neutral EU preferences with linear waiting cost Greedy is optimal in large markets.
  - Conjecture: holds for all smooth EU preferences (risk-averse/loving/non-linear waiting cost).

Robustness

- Single Type.
  - Simulations: results hold when there are not hard-to-match types $\lambda = -1$.
  - Much simpler single type model can not match the data and answer how hard-and-easy to match agents are differentially affected.
- Non-Vanishing Compatibility. (Ashlagi et al., 2013; Anderson et al., 2017; Akbarpour et al., 2020, 2019; Nikzad et al., 2019)
  - Simulations for $p = q = \frac{c}{\sqrt{m}}$: Greedy still outperforms Batching and Patient.
  - Simulations for $p = q = \frac{c}{m}$: Greedy has a lower waiting time and Patient a higher match rate.
- Different Objectives.
  - We show that for *all* risk-neutral EU preferences with linear waiting cost Greedy is optimal in large markets.
  - Conjecture: holds for all smooth EU preferences (risk-averse/loving/non-linear waiting cost).
  - Akbarpour et al., (2020) show that the loss ratio between different policies can be infinitely better under patient matching.

Conclusion

- We looked at **when to match** in large kidney exchanges.
- Compatibility graph:
  - No single-type model can match aggregate features of the data.
  - Simple and interpretable two type model that matches the data.
- Dynamic Matching:
  - Greedy matching is optimal in large markets.
    - For all risk-neutral EU preferences.
    - For hard and easy-to-match agents.
    - No trade-off between matching more agents and faster.
  - Empirically at the size of the NKR.
    - Greedy outperforms weekly, monthly, bimonthly matching.
    - Patient leads to $\equiv$ 1% higher match rate, but 35% longer waiting time.
    - Patient matching makes *easy-to-match* agents better of and hurts *hard-to-match* agents.

Thank You!

Notes

- "Small" markets:
  - Merging will increase the match rate (Agarwal et al. 2018, 2019). Emerging collaborations between European countries.
  - Chains will improve match rate and waiting times. Studies suggest that greedy does no harm (Anderson et al 2017, Ashlagi et al . 2017, Agarwal et al. 2018).
- How to match with heterogeneous match qualities? (the next talk...)