INTRODUCTION TO DATA SCIENCE

JOHN P DICKERSON

Lecture #27 – 5/4/2017

CMSC320
Tuesdays & Thursdays
3:30pm – 4:45pm

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

ANNOUNCEMENTS

HW4 is due today

- ~2/3 turned in already!
- Fill out the group assignment Google Drive document (see Piazza post for a private link)

Tuesday, May 9th will be a midterm & general wrap-up questions lecture day by Denis Peskov

- (Worked as a data scientist before coming to UMD.)

Any final tutorial questions ???????????

TODAY'S LECTURE

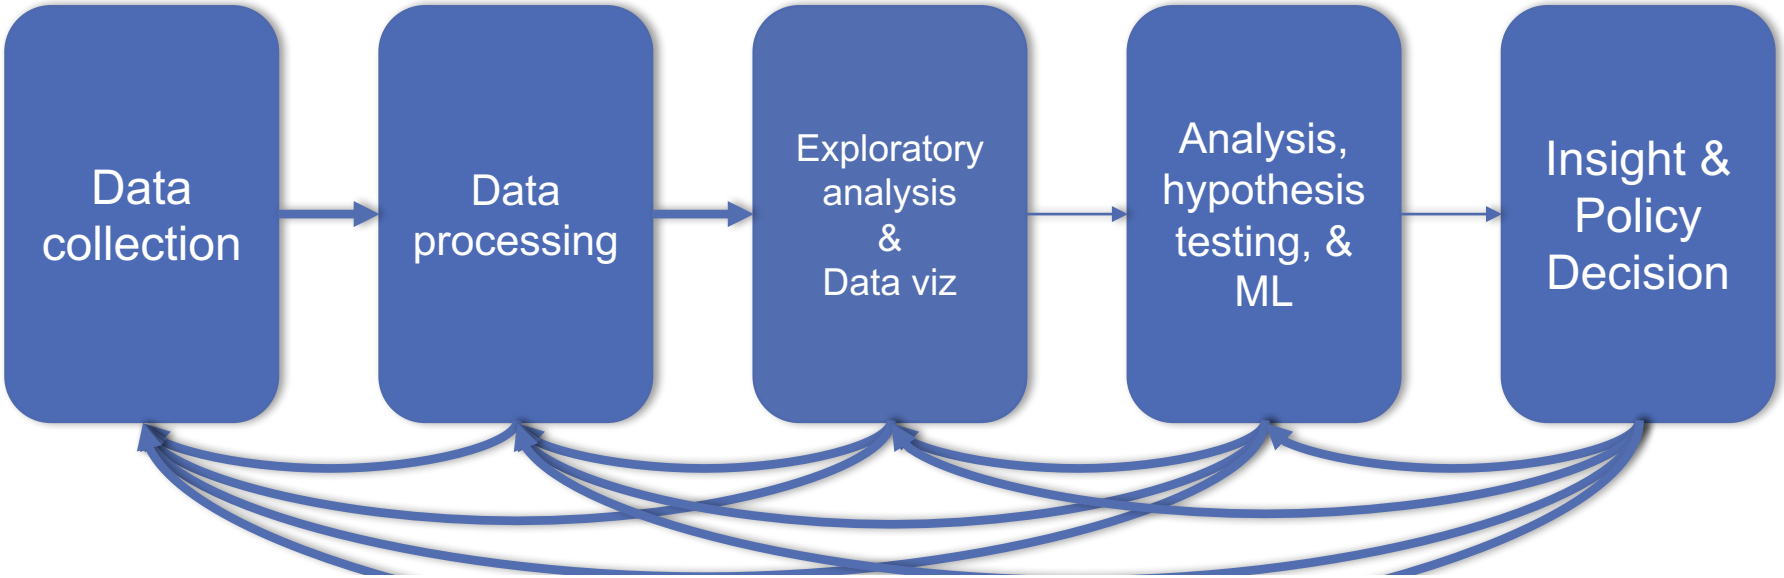

URGENT: Poll for Thursday's Lecture is now closed

A total of **40** vote(s) in **43** hours

20 (50% of users) Kidney Exchange

20 (50% of users) Neural Nets

Markets come in many forms …

… some of which don't conform to
conventional notions of markets …

… and some in which money may play little or no role.

– excerpt from *Who Gets What – and Why*

MATCHING MARKETS

In matching problems, prices do not do all – or any – of the work

Agents are paired with other (groups of) agents, transactions, or contracts

- Workers to firms
- Children to schools
- Residents to hospitals
- Patients to donors
- Advertisements to viewers
- Riders to rideshare drivers

UNCERTAINTY

- **Does a matched edge truly exist?**
- **How valuable is a match?**
- **Will a better match arrive in the future?**

upwork™
formerly oDesk

bing
淘宝网
Taobao.com

COMPETITION

Rival matching markets compete over the same agents

- **How does this affect global social welfare?**
- **How to differentiate?**

MATCH CADENCE

How quickly do new edges form?

How frequently does a market clear?

Is clearing centralized or decentralized?

Can agents reenter the market?

Use data & optimization – alongside human domain expertise – to learn matching policies

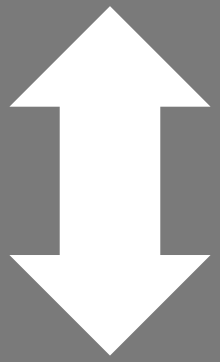

Strong theoretical underpinnings provide design guidance & runtime guarantees

THIS TALK

- **Three dimensions of matching market design:**
  - Managing short-term uncertainty
  - Balancing equity & efficiency
  - Combining human input and optimization
- ***(Each is supported by my work with a nationwide kidney exchange)***
- **Also, some open problems!**

Interested? CMSC498/CMSC499 Independent Research?
jpdickerson.com

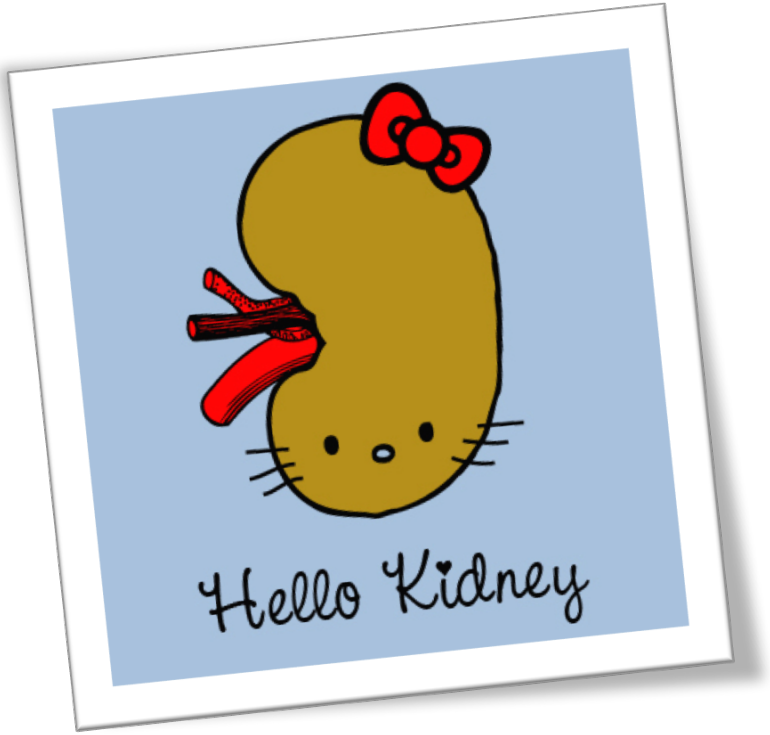

KIDNEY EXCHANGE

KIDNEY TRANSPLANTATION

- **US waitlist: about 100,000**
  - 36,157 added in 2015
- **4,553 people died while waiting**
- **12,236 people received a kidney from the deceased donor waitlist**

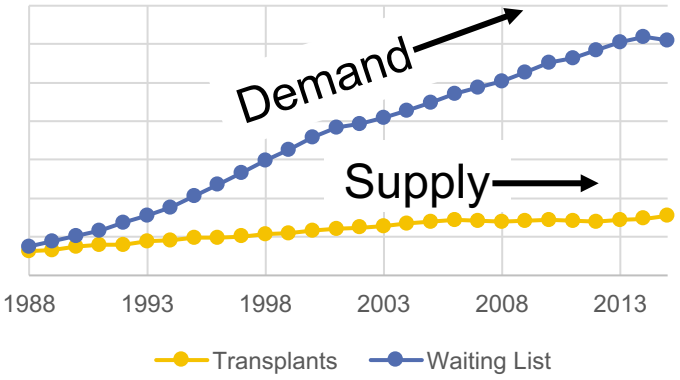

- **5,585 people received a kidney from a living donor**
  - Some through **kidney exchanges**! [Roth et al. 2004]
  - This talk: experience with UNOS national kidney exchange (and some data from the NHS NLDKSS)

UNOS DONATE LIFE

NHS

KIDNEY EXCHANGE

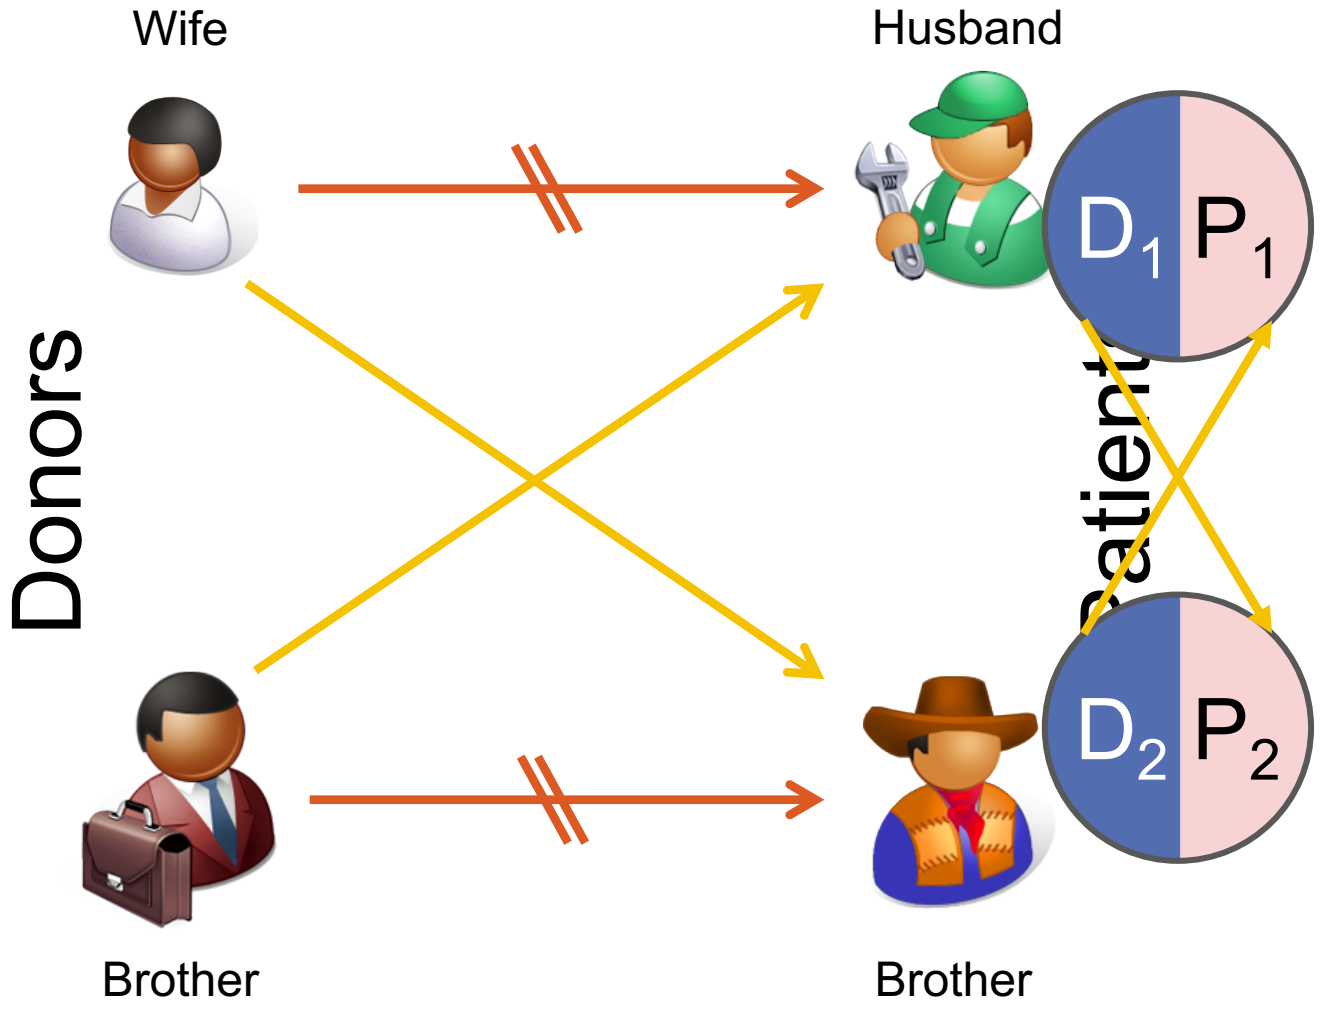

(2- and 3-cycles, all surgeries performed simultaneously)

NON-DIRECTED DONORS & CHAINS

[Rees et al. 2009]

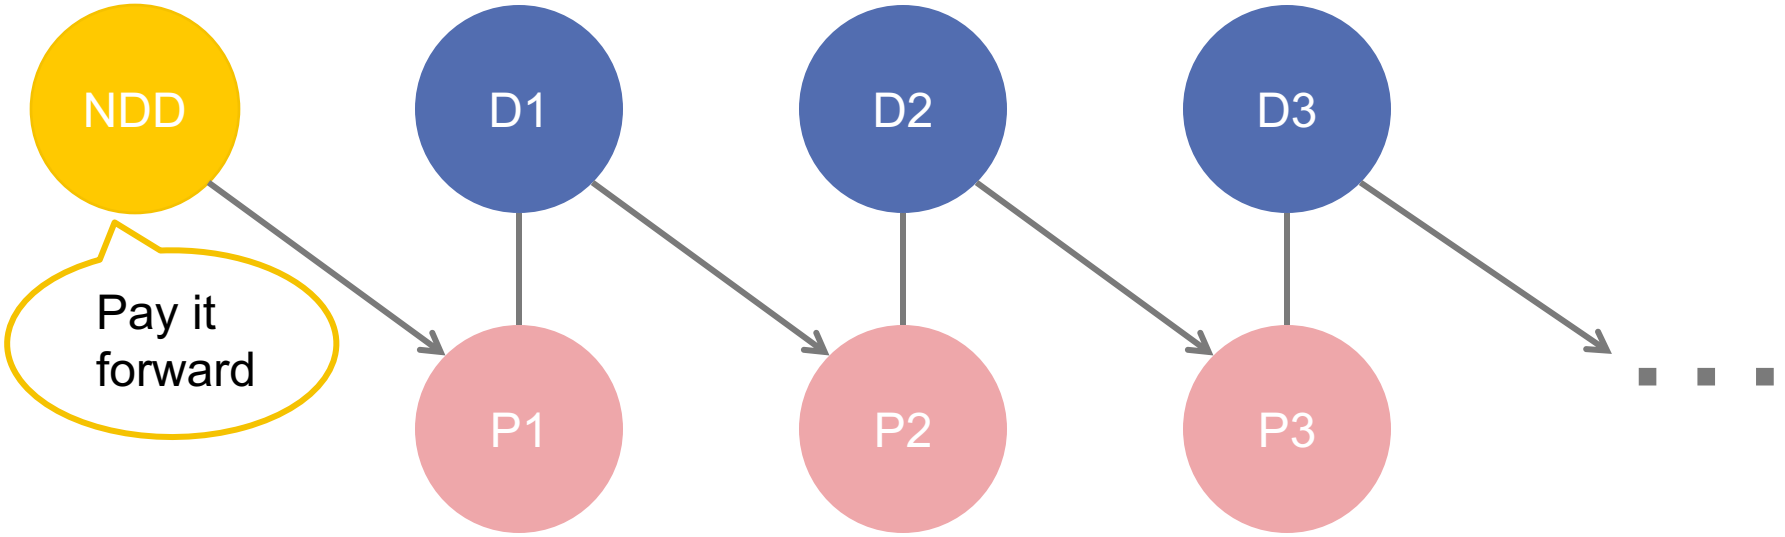

Not executed simultaneously, so no length cap required based on logistic concerns …

… **but in practice edges fail, so *some* finite cap is used!**

REAL-WORLD IMPACT

Kidney exchange is only a decade young, but already accounts for >10% of living donations in the United States

- Now a worldwide phenomenon (AU, CA, IL, PT, TR, UK, …)
- (Slowly) moving toward organized international exchange

Worked extensively with the United Network for Organ Sharing (UNOS) US nationwide kidney exchange!

- 153 transplant centers (roughly 66% of the US)
- Completely autonomous biweekly match runs
- Only automated exchange in the US

THE CLEARING PROBLEM

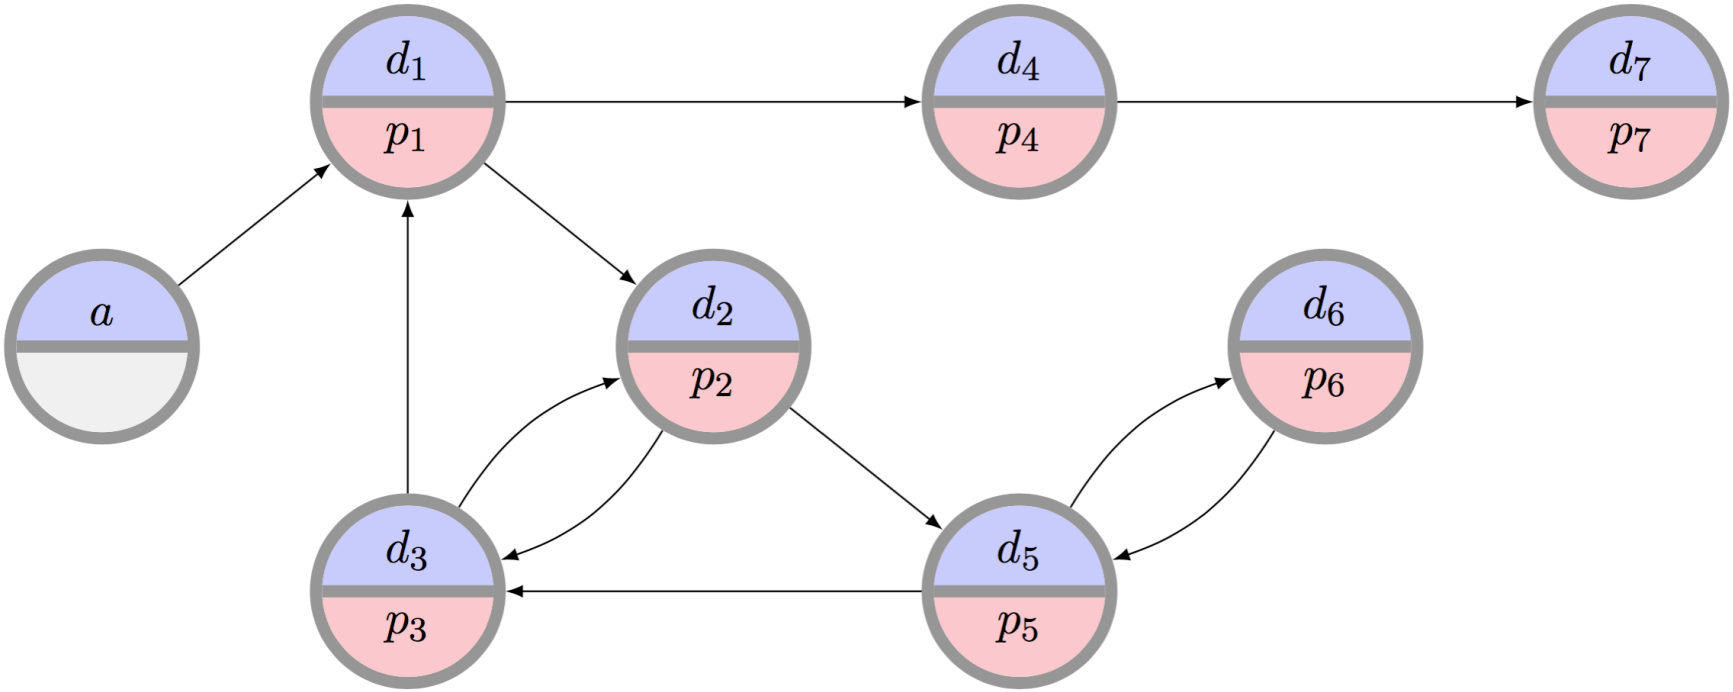

The clearing problem is to find the "best" disjoint set of cycles of length at most *L*, and chains

- Typically, $2 \leq L \leq 5$ for kidneys (e.g., L=3 at UNOS)
- NP-hard (for L>2) in theory, really hard in practice [Abraham et al. 07, Biro et al. 09] [Glorie et al. 2014, Anderson et al. 2015, Plaut et al. 2016, Dickerson et al. 2016 ...]

A SIMPLE INTEGER PROGRAM

("Best" = max weight, myopic matching)

[Roth et al. 04, 05, Abraham et al. 07]

Binary variable x_c for each feasible cycle or chain c

Maximize

$$u(M) = \Sigma\, w_c\, x_c$$

Subject to

$$\Sigma_{c\,:\,i \text{ in } c}\, x_c \leq 1 \text{ for each vertex } i$$

"SIMPLE" …?

THE BIG PROBLEM

What is "best"?

- Maximize matches right now or over time?
- Maximize transplants or matches?
- Prioritization schemes (i.e. fairness)?
- Modeling choices?
- Incentives? Ethics? Legality?

Optimization can handle this, but may be inflexible in hard-to-understand ways (for humans)

Want humans in the loop at a **high level**
(and then CS/Opt handles the implementation)

MANAGING SHORT-TERM UNCERTAINTY

[EC-13, EC-15, EC-16]

With A. Blum, N. Haghtalab, D. Manlove, B. Plaut, A. Procaccia, T. Sandholm, A. Sharma, J. Trimble

MATCHED ≠ TRANSPLANTED

Only around 10-15% of UNOS matched structures result in an actual transplant

- Similarly low % in other exchanges [ATC 2013]

***Many* reasons for this. How to handle?**

One way: encode *probability of transplantation* rather than just feasibility

- for individuals, cycles, chains, and full matchings

FAILURE-AWARE MODEL

Compatibility graph *G*

- Edge (v_i, v_j) if v_i's donor can donate to v_j's patient
- Weight w_e on each edge e

Success probability q_e for each edge e

Discounted utility of cycle c

$$u(c) = \sum w_e \cdot \prod q_e$$

Value of successful cycle

Probability of success

FAILURE-AWARE MODEL

Discounted utility of a *k*-chain *c*

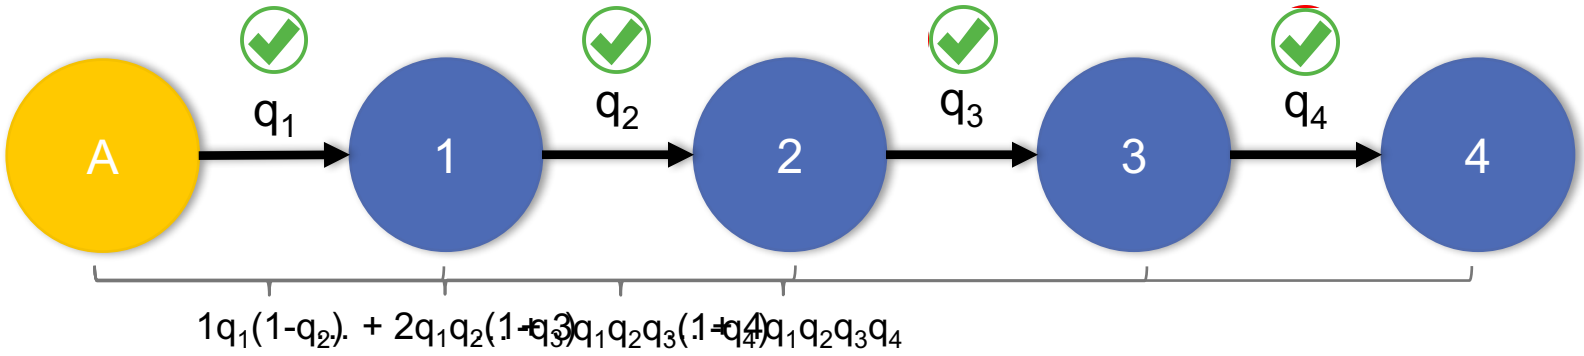

$$u(c) = \left[\sum_{i=1}^{k-1}(1-q_i)\, i \prod_{j=0}^{i-1} q_j\right] + \left[k \prod_{i=0}^{k-1} q_i\right]$$

Exactly first *i* transplants

Chain executes in entirety

Cannot simply "reweight by failure probability"

DISCOUNTED CLEARING PROBLEM

("Best" = max expected cardinality | limited recourse)

Find matching M^* with highest discounted utility

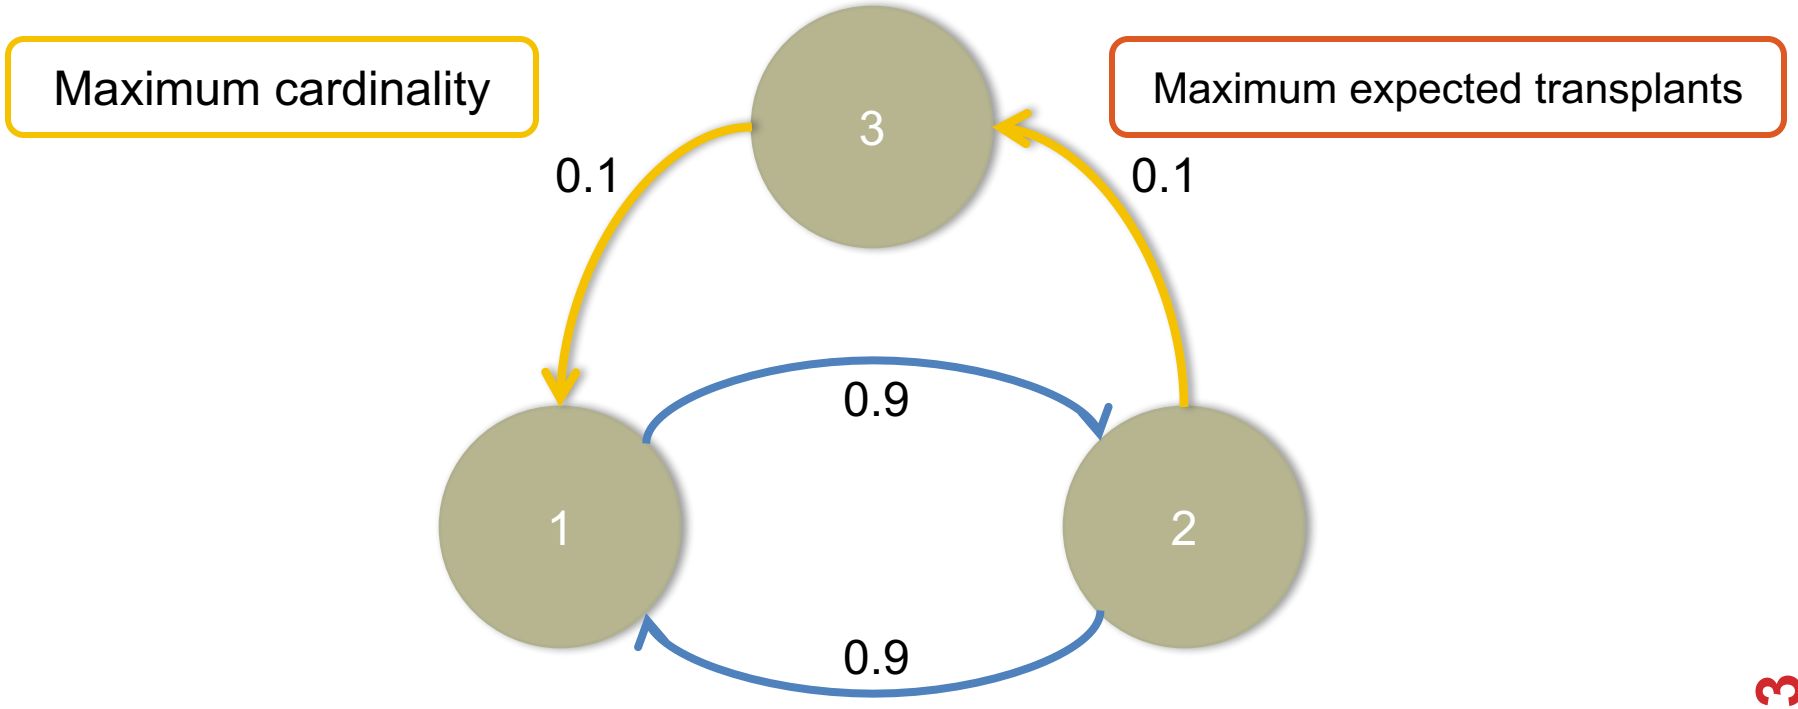

“IN THE LARGE”

G(n, t(n), p): random graph with

- n patient-donor pairs
- t(n) altruistic donors
- Probability Θ(1/n) of incoming edges

Constant transplant success probability q

THEOREM

For all $q \in (0,1)$ and $\alpha, \beta > 0$, given a large $G(n, \alpha n, \beta/n)$, w.h.p. there exists some matching M' s.t. for every maximum cardinality matching M,

$$u_q(M') \geq u_q(M) + \Omega(n)$$

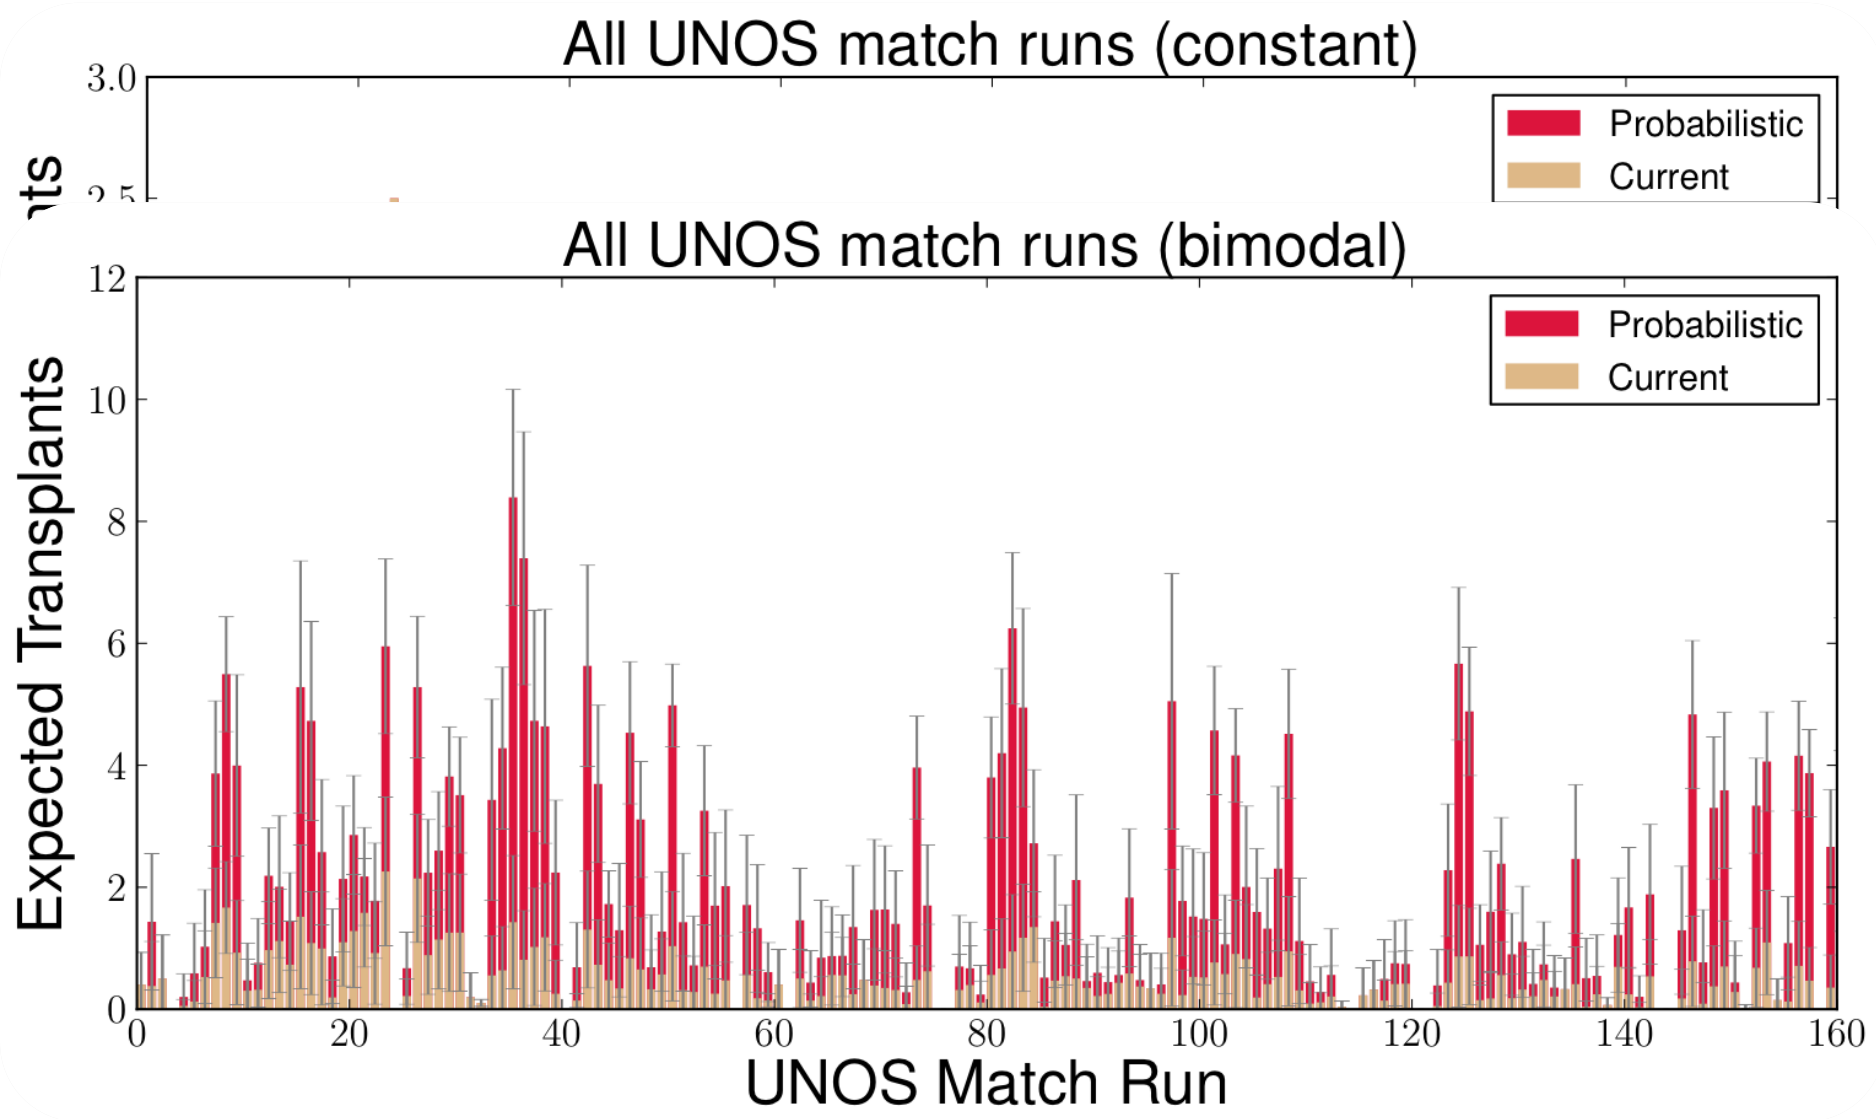

Under discussion for implementation at UNOS

PRE-MATCH EDGE TESTING

Idea: perform a small amount of costly testing before a match run to test for (non)existence of edges

- E.g., more extensive medical testing, donor interviews, surgeon interviews, …

Cast as a stochastic matching (or set packing) problem:

Given a graph $G(V,E)$, choose subset of edges S such that:

$$|M(S)| \geq (1-\varepsilon)\,|M(E)|$$

Need: "sparse" S, where every vertex has O(1) incident tested edges

Select edges in rounds of at most 1 edge incident per vertex

ADAPTIVE ALGORITHM

For *R* rounds, do:

1. Pick a max-cardinality matching *M* in graph *G*, minus already-queried edges that do not exist
2. Query all edges in *M*

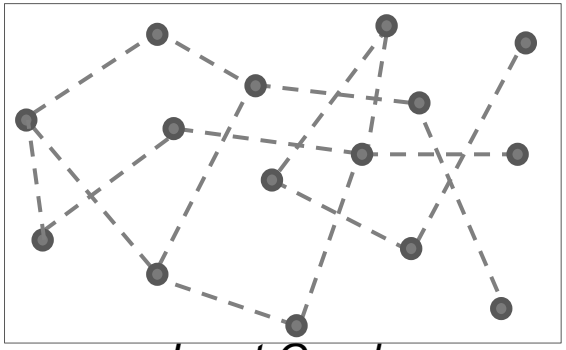

Input Graph

r	Base graph	Matching picked	Result of queries
1:	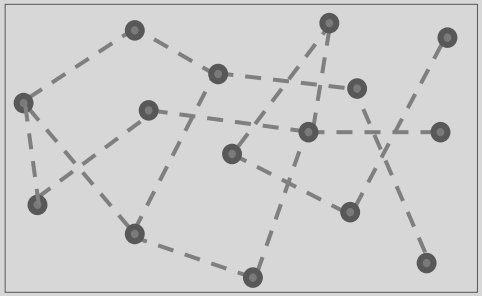	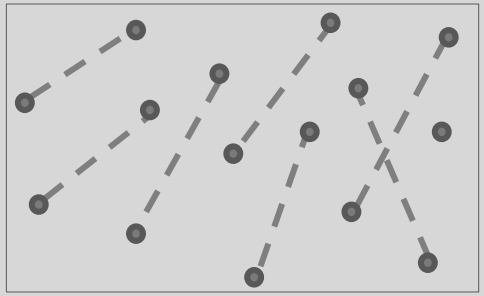	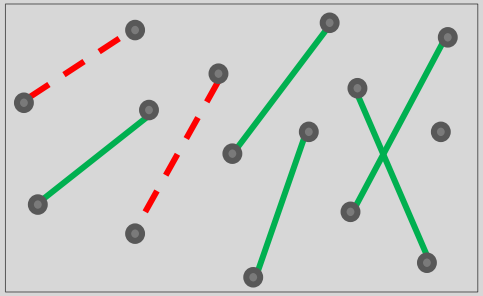
2:	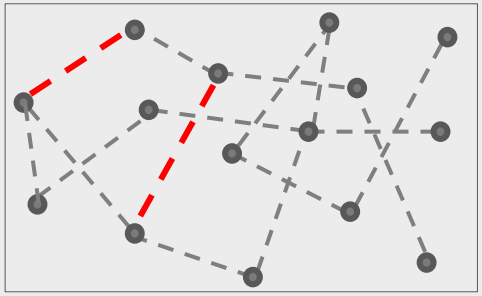	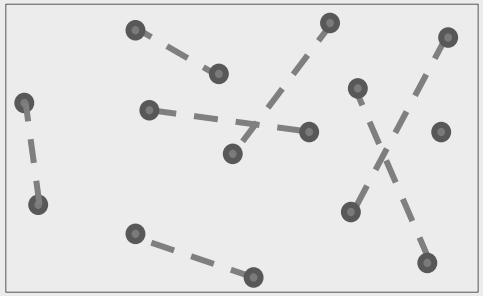	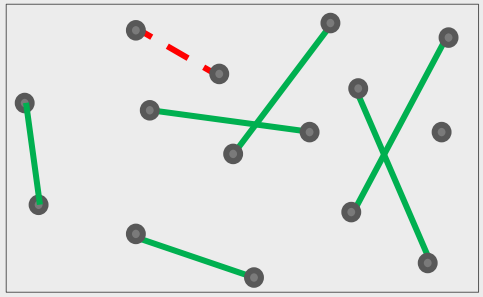

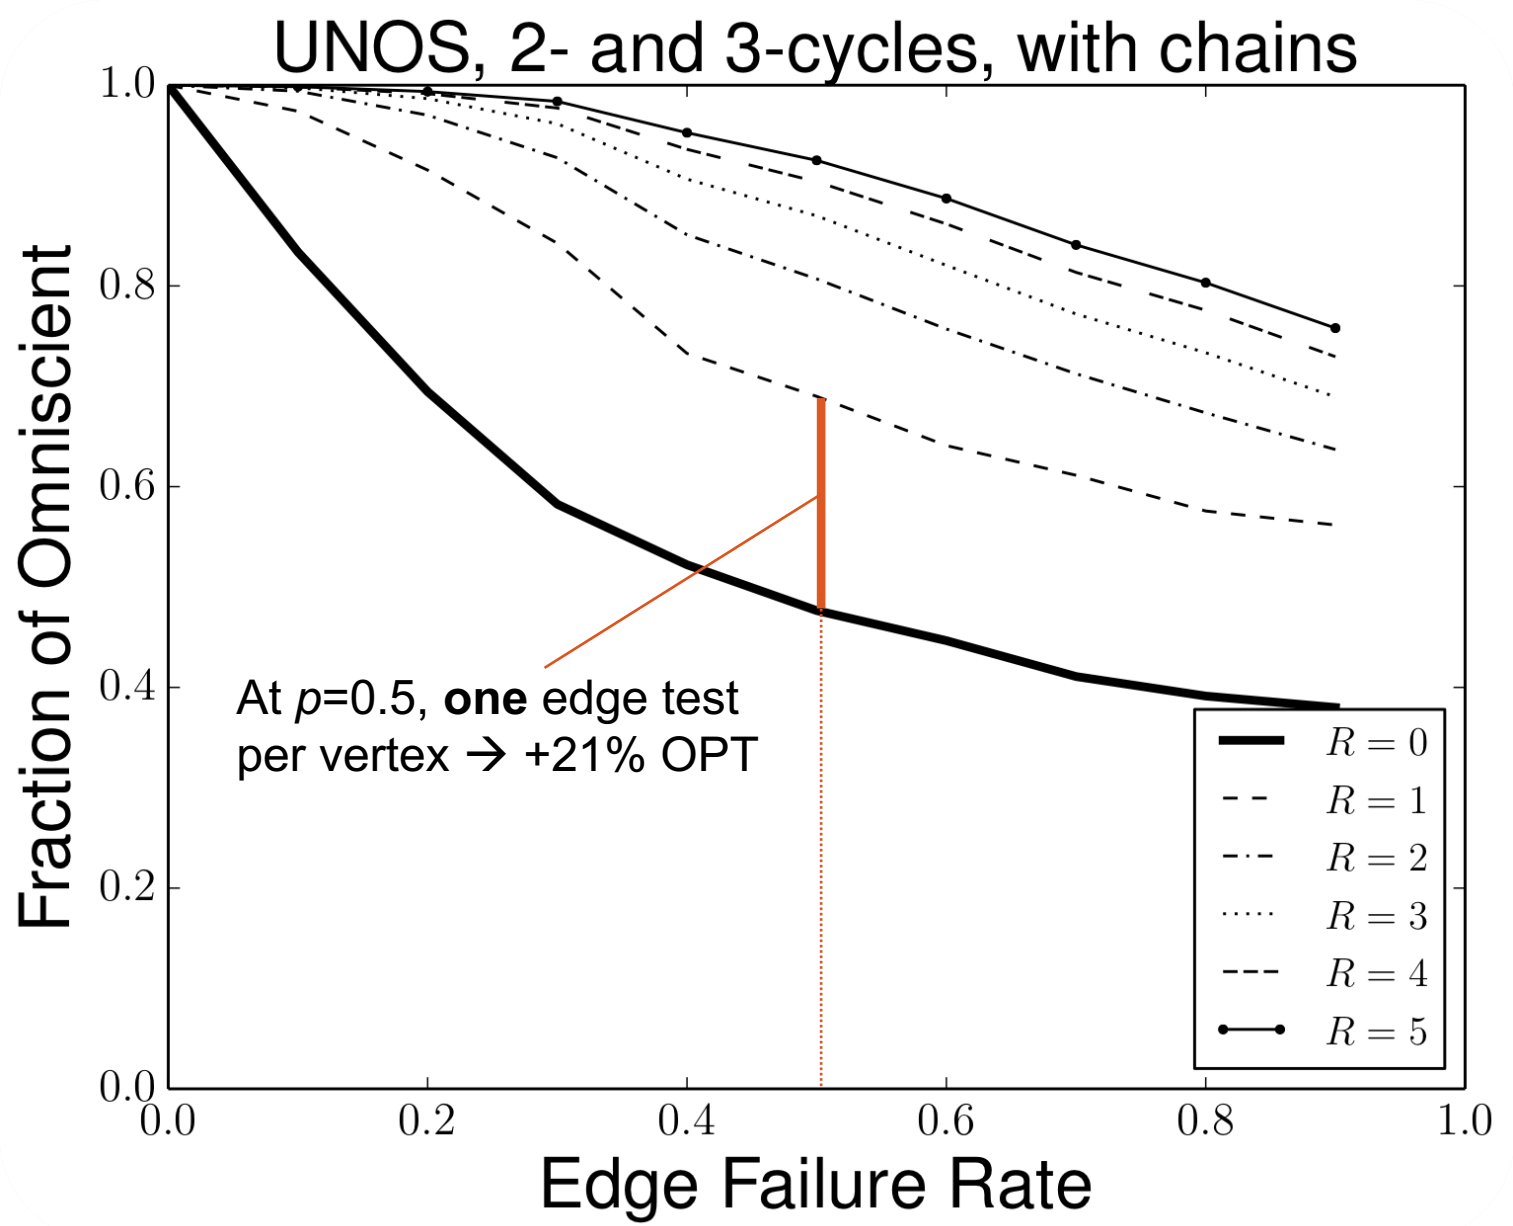

Even 1 or 2 extra tests would result in a huge lift

In theory and practice, we're helping the global bottom line by considering post-match failure …

… But can this hurt some individuals?

BALANCING EQUITY AND EFFICIENCY

[AAMAS-14, AAAI-15, Invited to AIJ]

With A. Procaccia and T. Sandholm

SENSITIZATION AT UNOS

Highly-sensitized patients: unlikely to be compatible with a random donor

- Deceased donor waitlist: 17%
- Kidney exchanges: **much** higher (60%+)

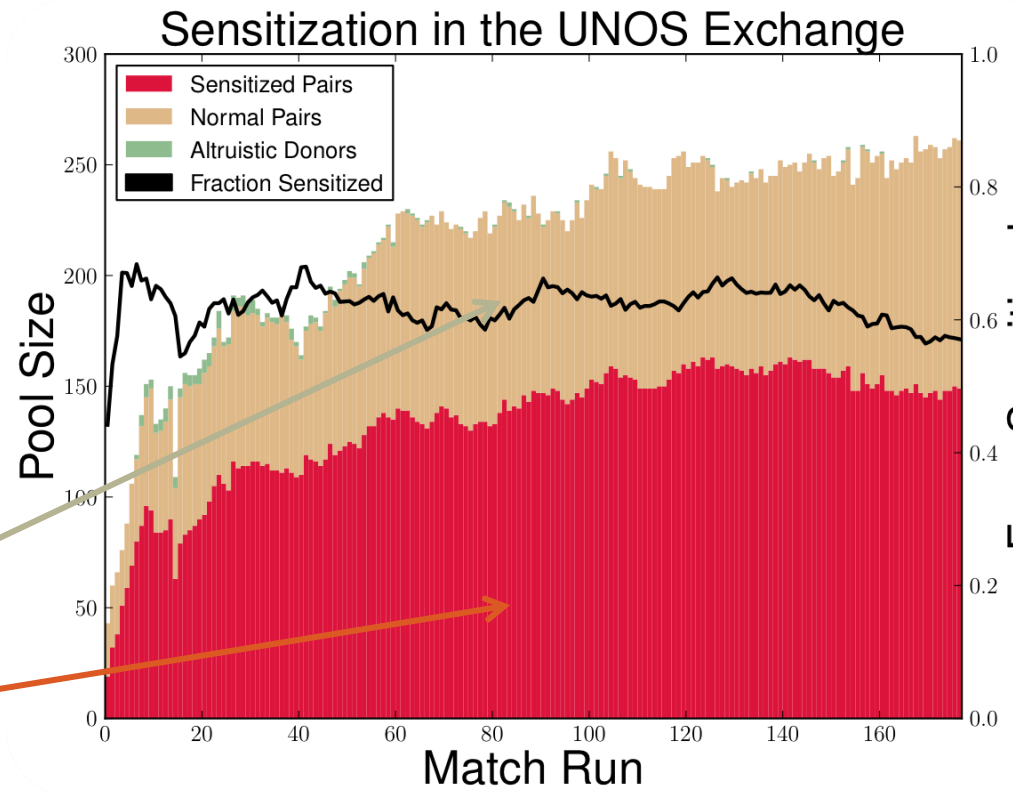

"Easy to match" patients

"Hard to match" patients

PRICE OF FAIRNESS

Efficiency vs. fairness:

- **Utilitarian** objectives may favor certain classes at the expense of marginalizing others
- **Fair** objectives may sacrifice efficiency in the name of egalitarianism

Price of fairness: relative system efficiency loss under a fair allocation [Bertismas, Farias, Trichakis 2011] [Caragiannis et al. 2009]

PRICE OF FAIRNESS IN KIDNEY EXCHANGE

Want a matching M^* that maximizes utility function $u: \mathcal{M} \to \mathbb{R}$

$$M^* = \operatorname*{argmax}_{M \in \mathcal{M}} u(M)$$

Price of fairness: relative loss of match **efficiency** due to **fair** utility function u_f:

$$POF(\mathcal{M}, u_f) = \frac{u(M^*) - u(M_f^*)}{u(M^*)}$$

TOWARD USABLE FAIRNESS RULES

In healthcare, important to work within (or near to) the constraints of the fielded system

- [Bertsimas, Farias, Trichakis 2013]
- Our experience with UNOS

This talk: simplest and most intuitive rule:

- **Weighted:** implementation of “priority points”

WEIGHTED FAIRNESS

Value matching a highly-sensitized patient at (1+β) that of a lowly-sensitized patient, β>0

Re-weighting is a preprocess →

works with all present-day exchange solvers

Difficult to find a "good" β?

- Empirical exploration helps strike a balance

CONTRADICTORY GOALS

Earlier, we saw failure-aware matching results in tremendous gains in #expected transplants

Gain comes at a price – may further marginalize hard-to-match patients because:

- Highly-sensitized patients → matched in chains
- Highly-sensitized patients → higher failure rates

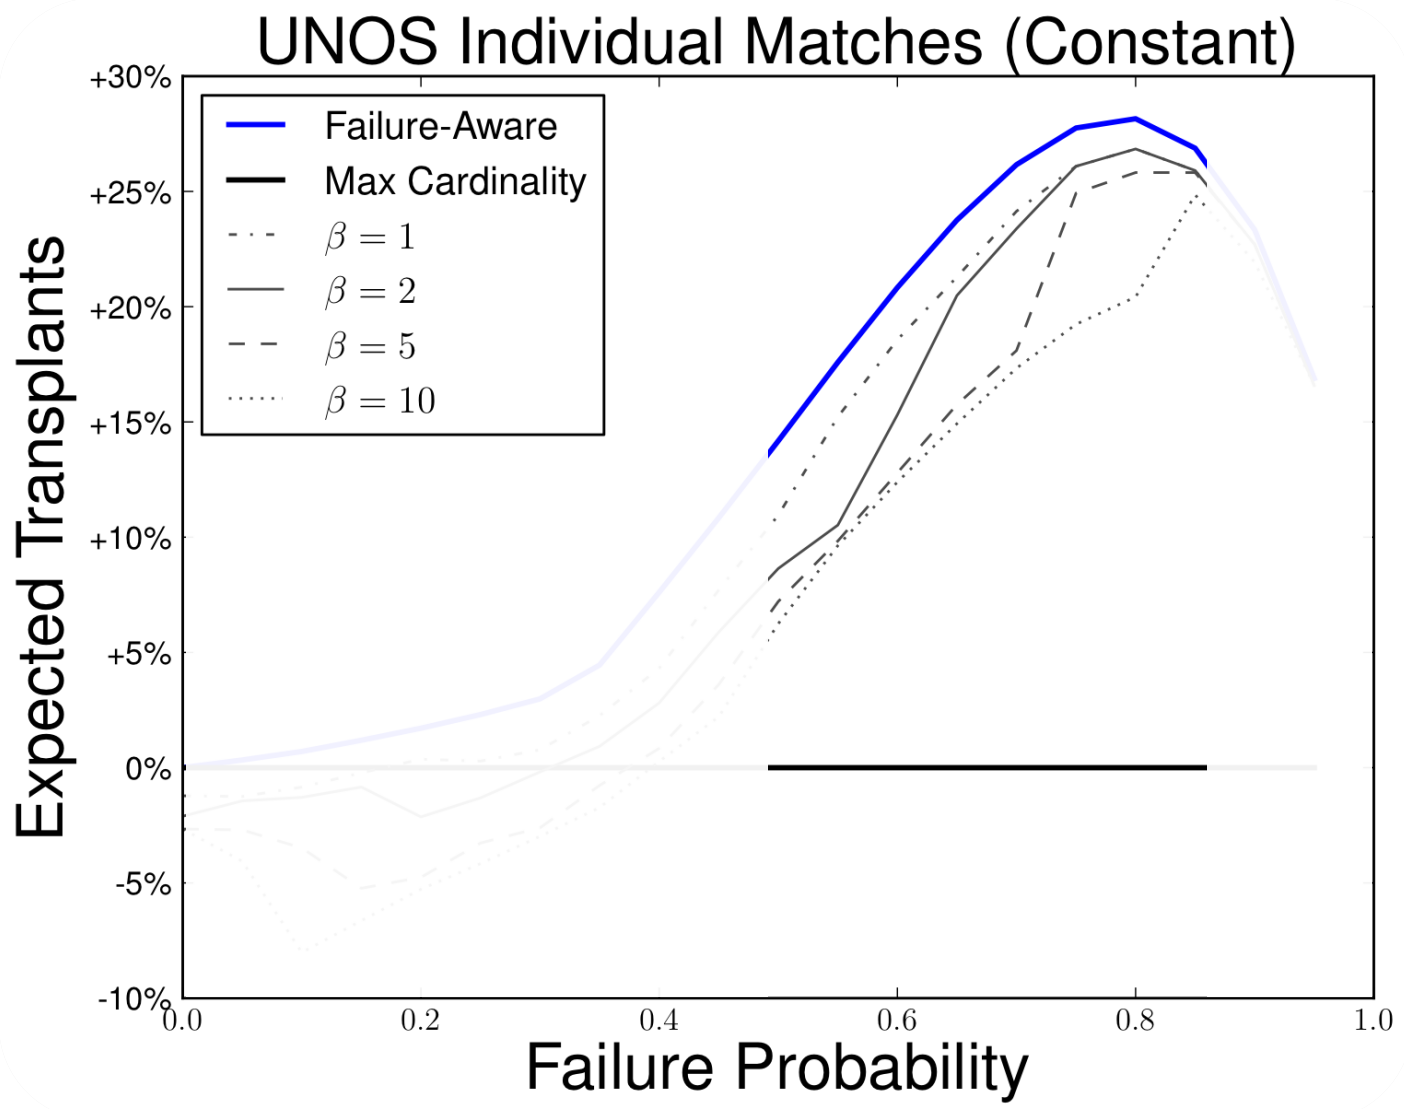

UNOS runs, weighted fairness, constant probability of failure (x-axis), increase in expected transplants over deterministic matching (y-axis)

Fairness vs. efficiency can be balanced in theory and in practice in a static model …

… But how should we match over time?

LEARNING TO MATCH IN A DYNAMIC ENVIRONMENT

[AAAI-12, AAAI-15, NIPS-15 MLHC]

With A. Procaccia and T. Sandholm

DYNAMIC KIDNEY EXCHANGE

Kidney exchange is a naturally dynamic event

Can be described by the evolution of its graph

- Additions, removals of edges and vertices

Vertex Removal	Edge Removal	Vertex/Edge Add
Transplant, this exchange	Matched, positive crossmatch	Normal entrance
Transplant, deceased donor waitlist	Matched, candidate refuses donor	
Transplant, other exchange ("sniped")	Matched, donor refuses candidate	
Death or illness	Pregnancy, sickness changes HLA	
Altruist runs out of patience		
Bridge donor reneges		

How should we balance matching now versus waiting to match?

FUTUREMATCH: LEARNING TO MATCH IN DYNAMIC ENVIRONMENTS

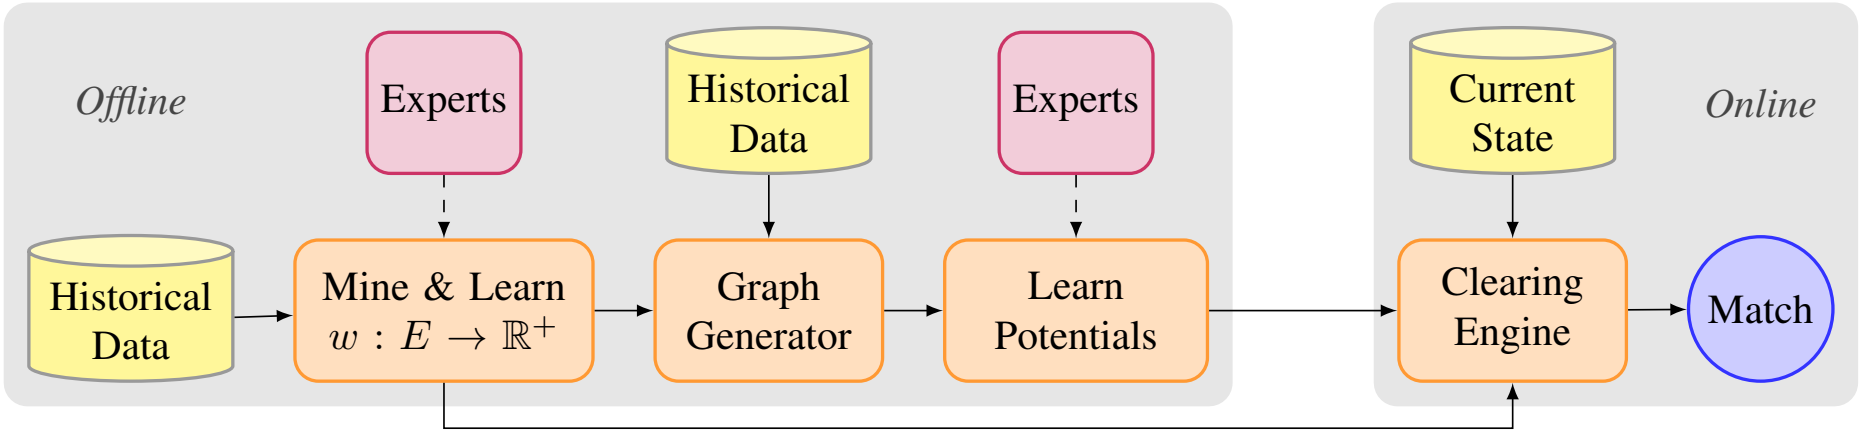

Offline (run once or periodically)

1. Domain expert describes overall goal
2. Take historical data and policy input to learn a weight function *w* for match quality
3. Take historical data and create a graph generator with edge weights set by *w*
4. Using this generator and a realistic exchange simulator, learn potentials for graph elements as a function of the exchange dynamics

Online (run every match)

1. Combine *w* and potentials to form new edge weights on real input graphs
2. Solve maximum weighted matching and return match

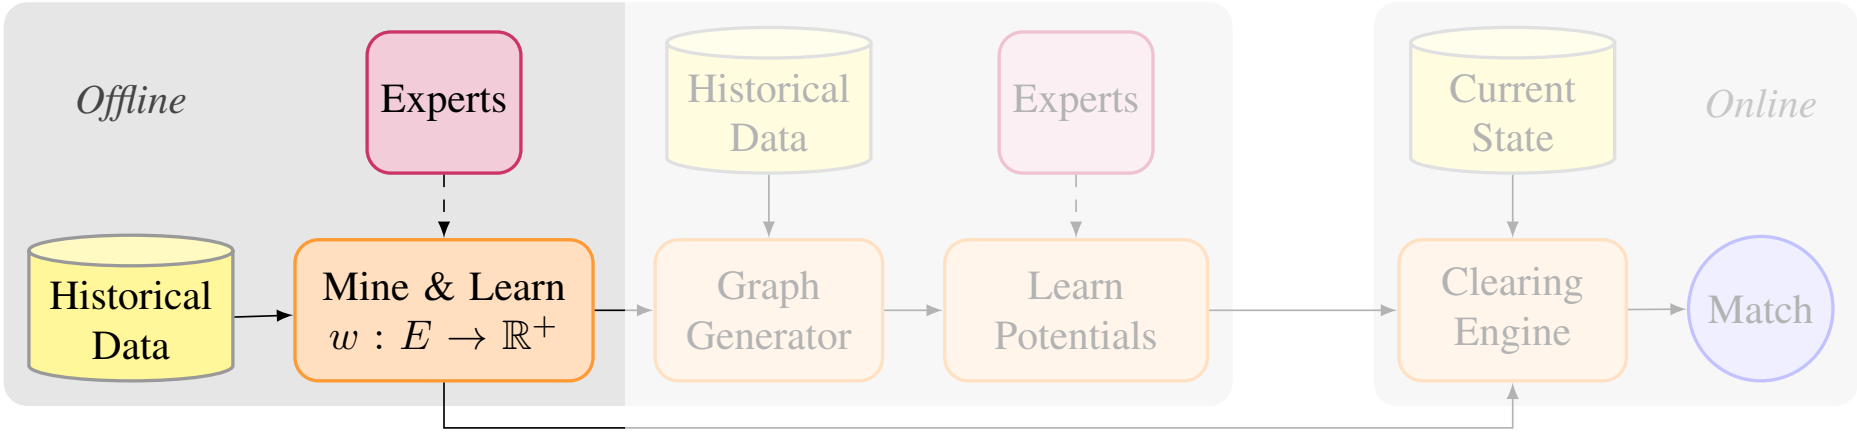

Example objective (MaxLife)

- Maximize aggregate length of time donor organs last in patients …
  - … possibly subject to prioritization schemes, fairness, etc …
- Learn survival rates from all living donations since 1987
  - ~75,000 transplants
- Translate to edge weight

Imperfect HLA match
has worse survival rate than
perfect HLA match

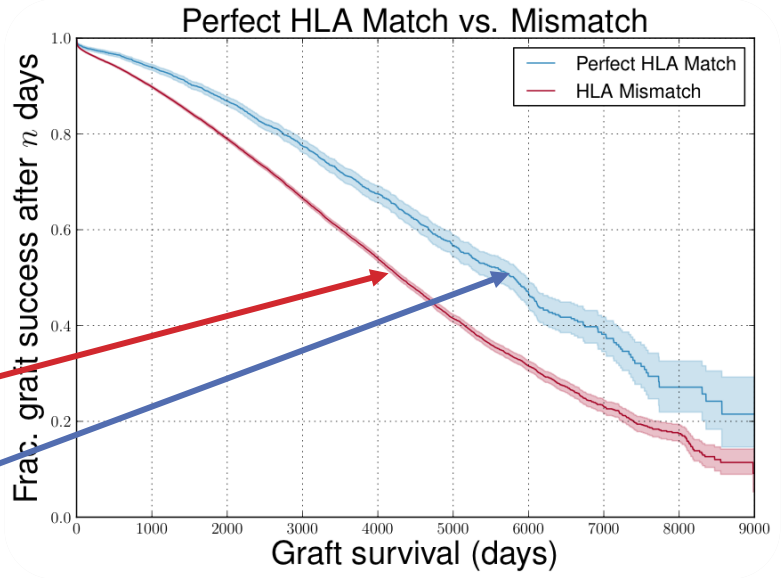

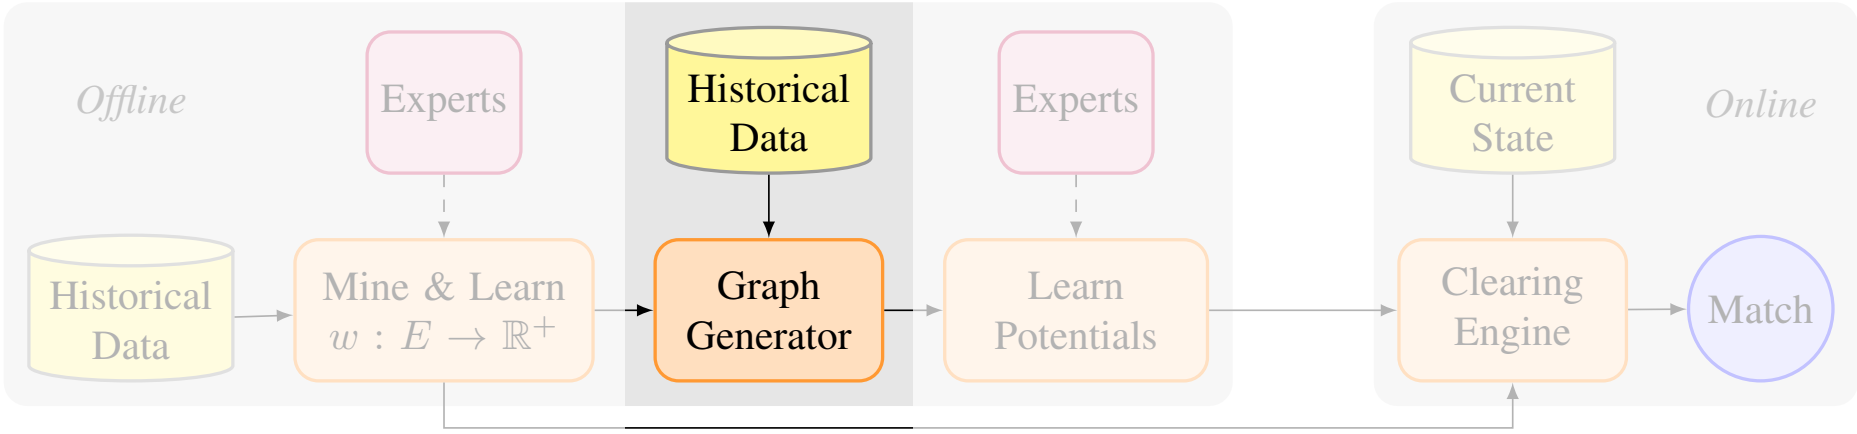

300+ match runs with real UNOS data

Important to use realistic distribution

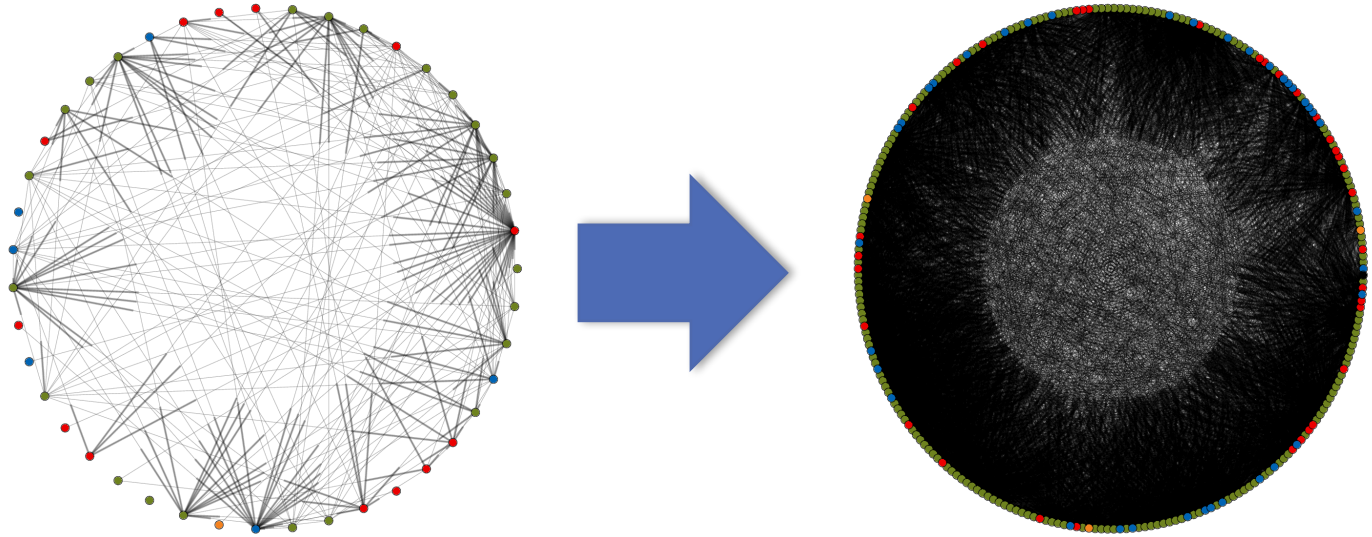

UNOS
(first match run)

UNOS
(recent snapshot)

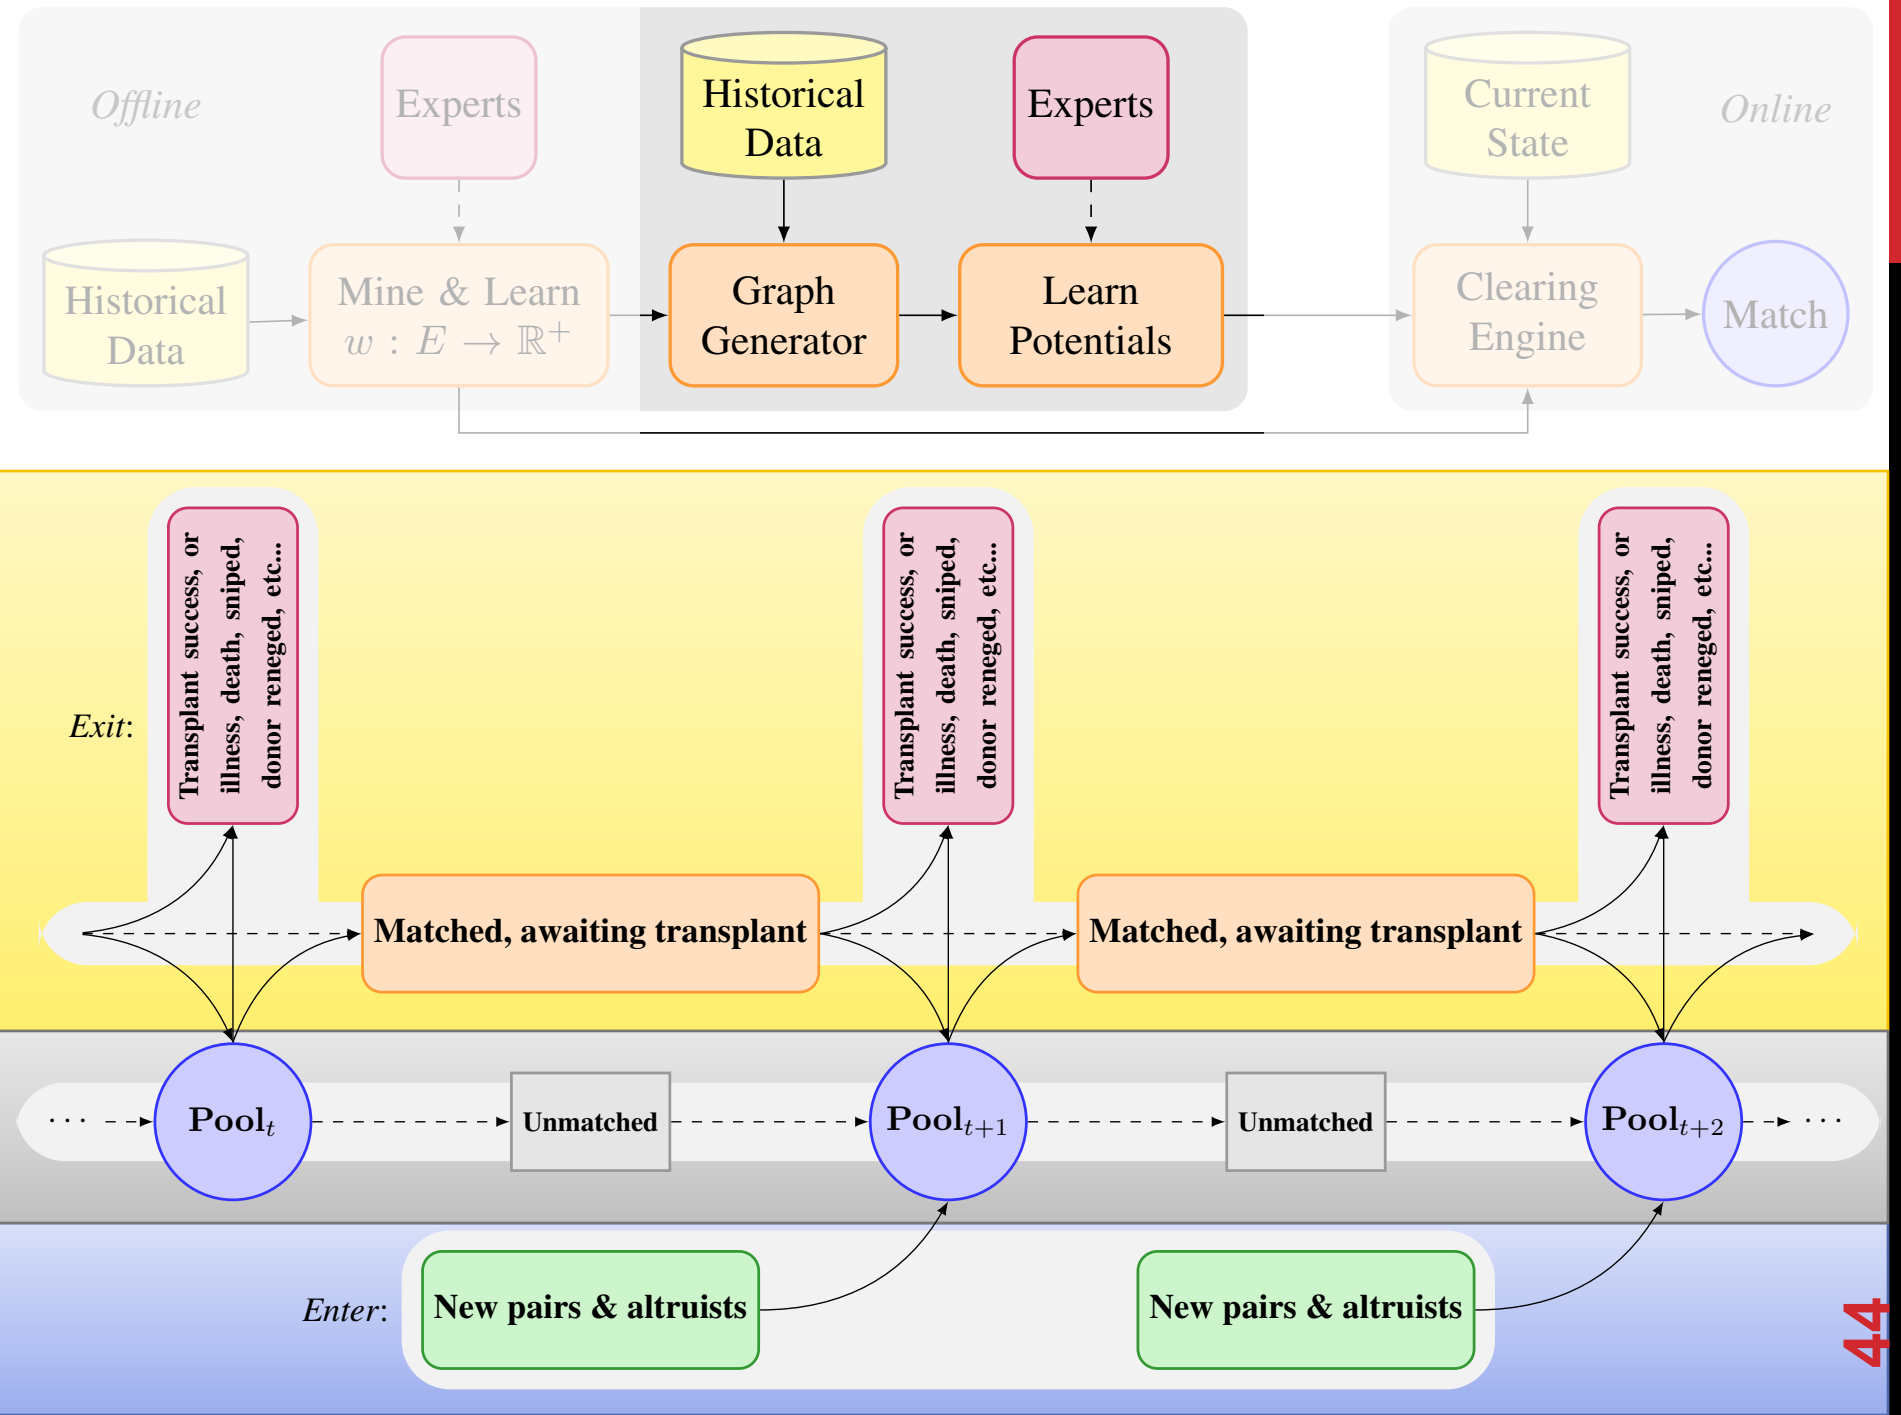
Offline
Experts
Historical Data
Experts
Current State
Online
Historical Data
Mine & Learn
$w : E \to \mathbb{R}^+$
Graph Generator
Learn Potentials
Clearing Engine
Match
Exit:
Transplant success, or illness, death, sniped, donor reneged, etc...
Transplant success, or illness, death, sniped, donor reneged, etc...
Transplant success, or illness, death, sniped, donor reneged, etc...
Matched, awaiting transplant
Matched, awaiting transplant
$\cdots$
Pool_t
Unmatched
Pool_{t+1}
Unmatched
Pool_{t+2}
$\cdots$
Enter:
New pairs & altruists
New pairs & altruists

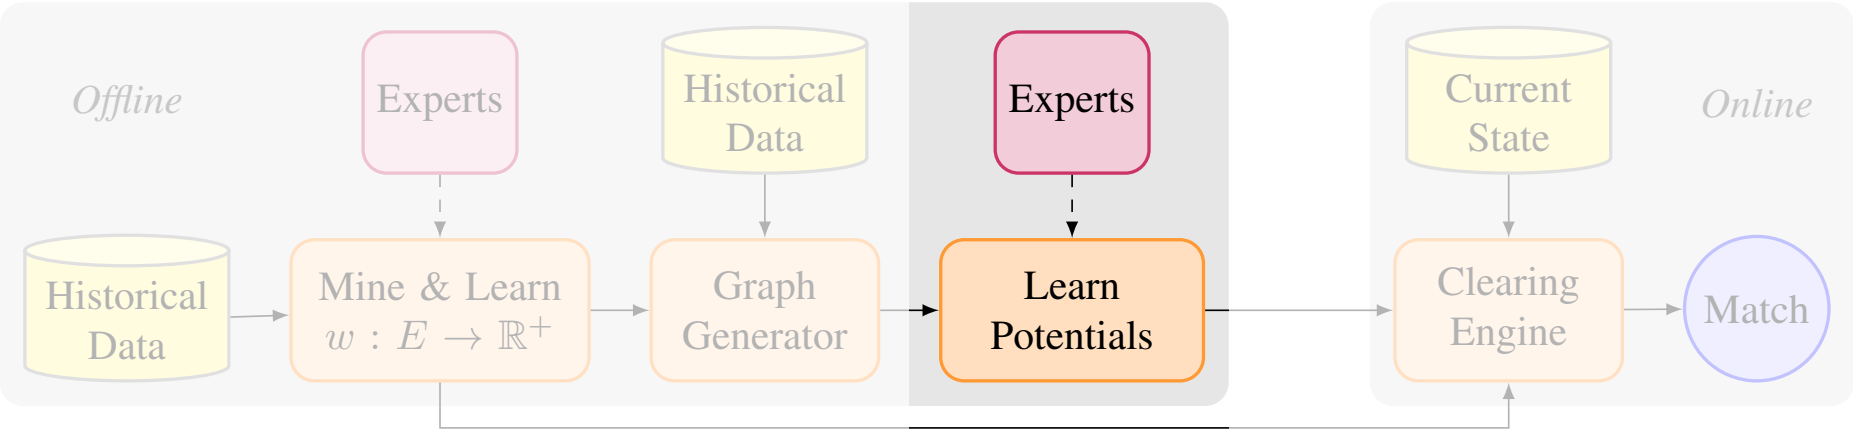

Full optimization problem is very difficult

- Realistic theory is too complex
- Trajectory-based methods do not scale

Approximation idea:

- Associate with each "element type" its **potential** to help objective in the future
- (Must learn these potentials)
- Combine potentials with edge weights, perform myopic maximum utility matching

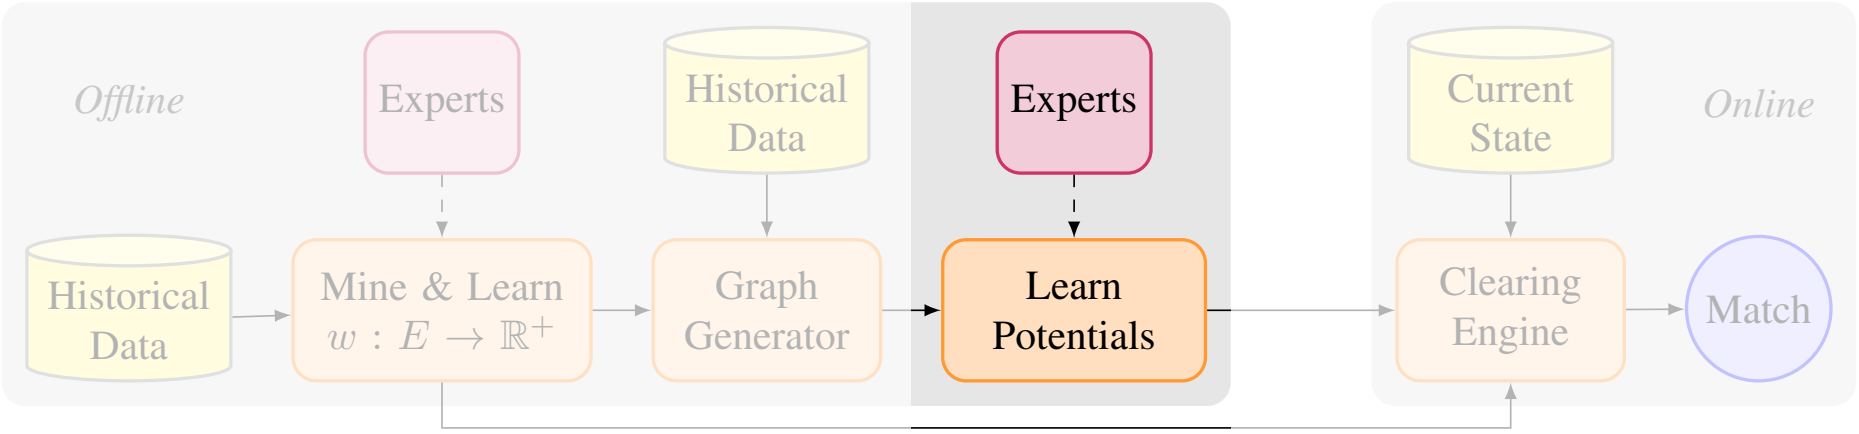

What is a potential?

Given a set of features *Θ* representing structural elements (e.g., vertex, edge, subgraph type) of a problem:

- The potential P_θ for a type θ quantifies the **future usefulness** of that element

E.g., let *Θ* = {O-O, O-A, …, AB-AB, •-O, …, •-AB}

- 16 patient-donor types, 4 altruist types
- O-donors better than A-donors, so: $P_{\bullet\text{-}O} > P_{\bullet\text{-}A}$

Heavy one-time computation to learn potential of each type *θ* – we use SMAC

[Hutter Hoos Leyton-Brown 2011]

PITTSBURGH SUPERCOMPUTING CENTER

IBM

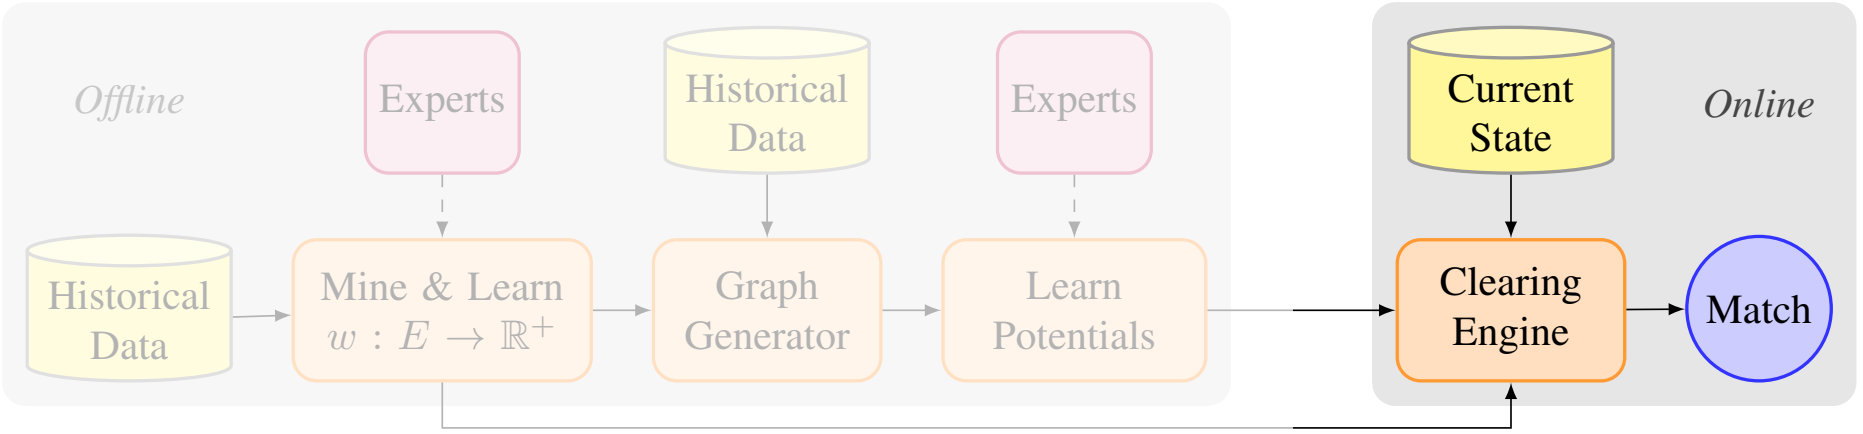

Online:

Adjust solver to take potentials into account at runtime

- E.g., $P_{\bullet\text{-}O} = 2.1$ and $P_{O\text{-}AB} = 0.1$
- Edges between O-altruist and O-AB pair has weight:
  $1 - 0.5(2.1+0.1) = -0.1$
- Chain must be long enough to offset negative weight

Also take into account learned weight function *w*

Edge weight preprocess → no or low runtime hit!

EXPERIMENTAL RESULTS & IMPACT

PITTSBURGH SUPERCOMPUTING CENTER

Presented at Supercomputing-14
Tied with IBM Watson

We show it is possible to:

- Increase overall #transplants a lot at a (much) smaller decrease in #marginalized transplants
- Increase #marginalized transplants a lot at no or very low decrease in overall #transplants
- Increase both #transplants and #marginalized

Sweet spot depends on distribution:

- Luckily, we can generate – and learn from – realistic families of graphs!

FutureMatch has been used for policy recommendations at UNOS

THE CUTTING EDGE

MANAGING INCENTIVES

Clearinghouse cares about global welfare:

- How many patients received kidneys (over time)?

Transplant centers care about their individual welfare:

- How many of my own patients received kidneys?

Patient-donor pairs care about their individual welfare:

- Did I receive a kidney?
- (Most work considers just clearinghouse and centers)

PRIVATE VS GLOBAL MATCHING

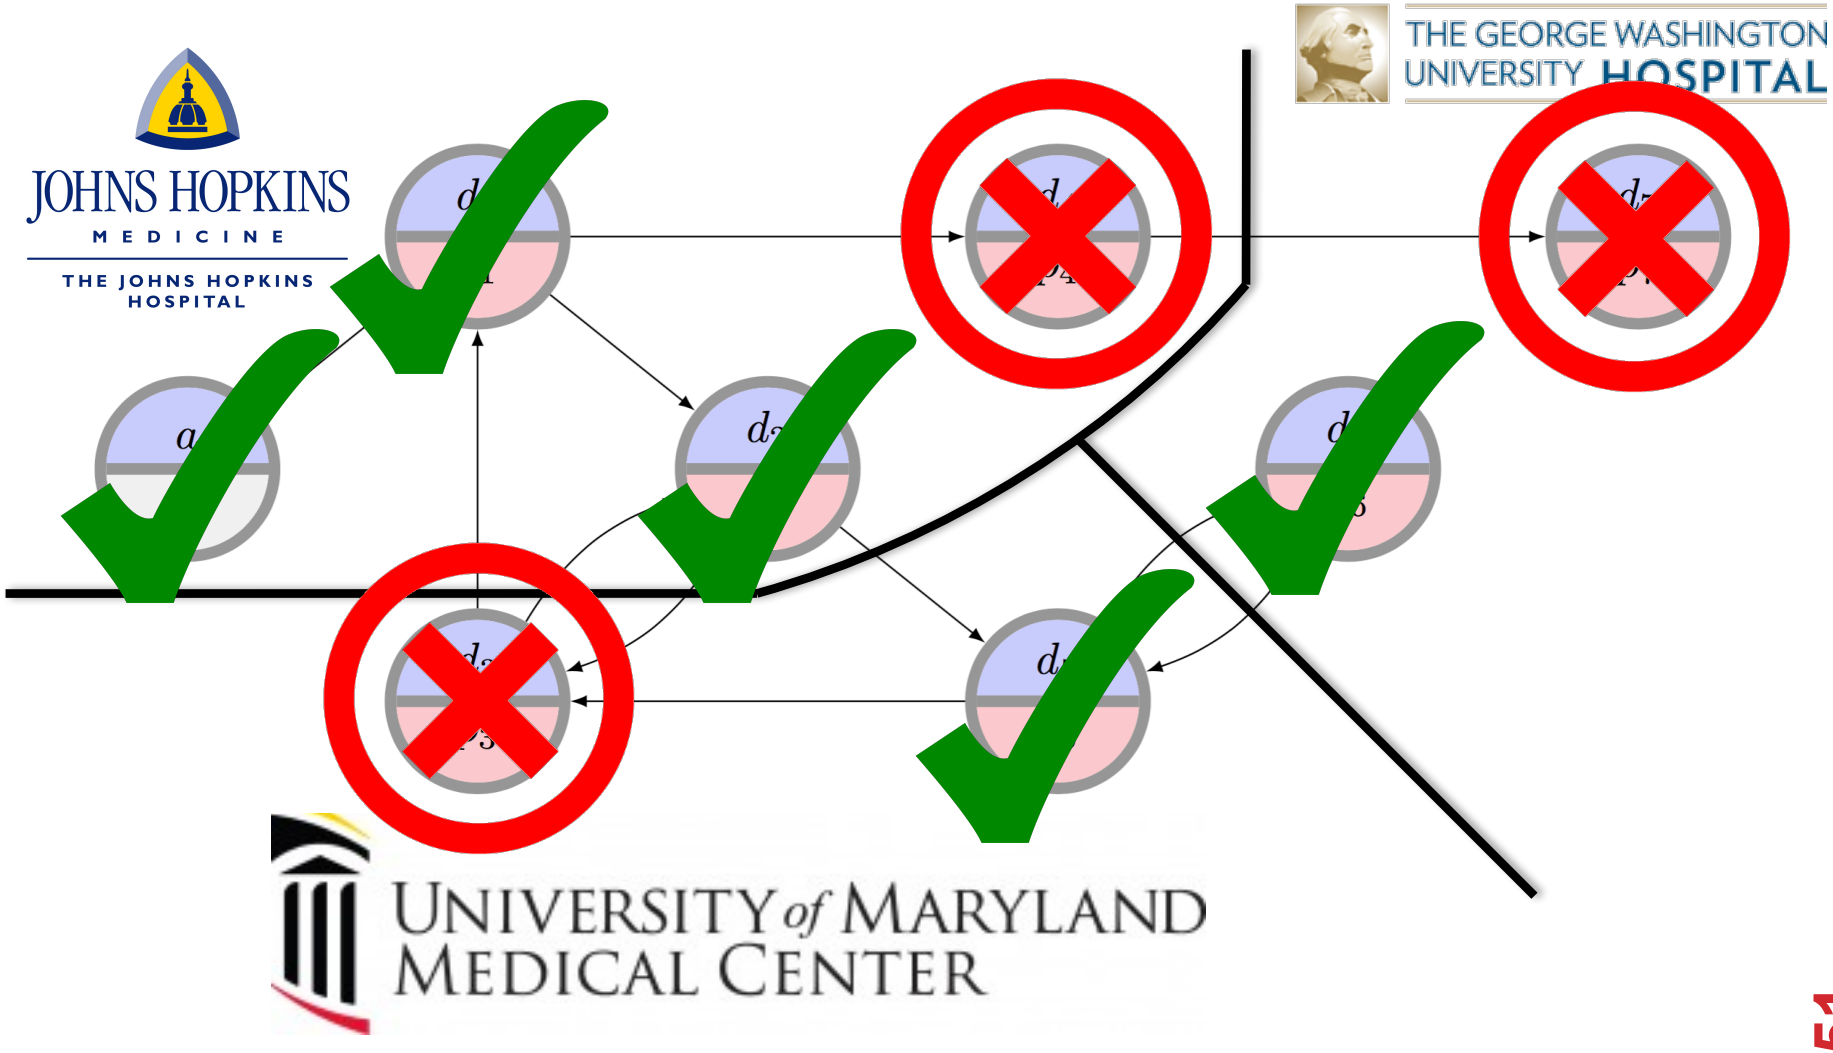

(Example transplant centers; not real!)

0%

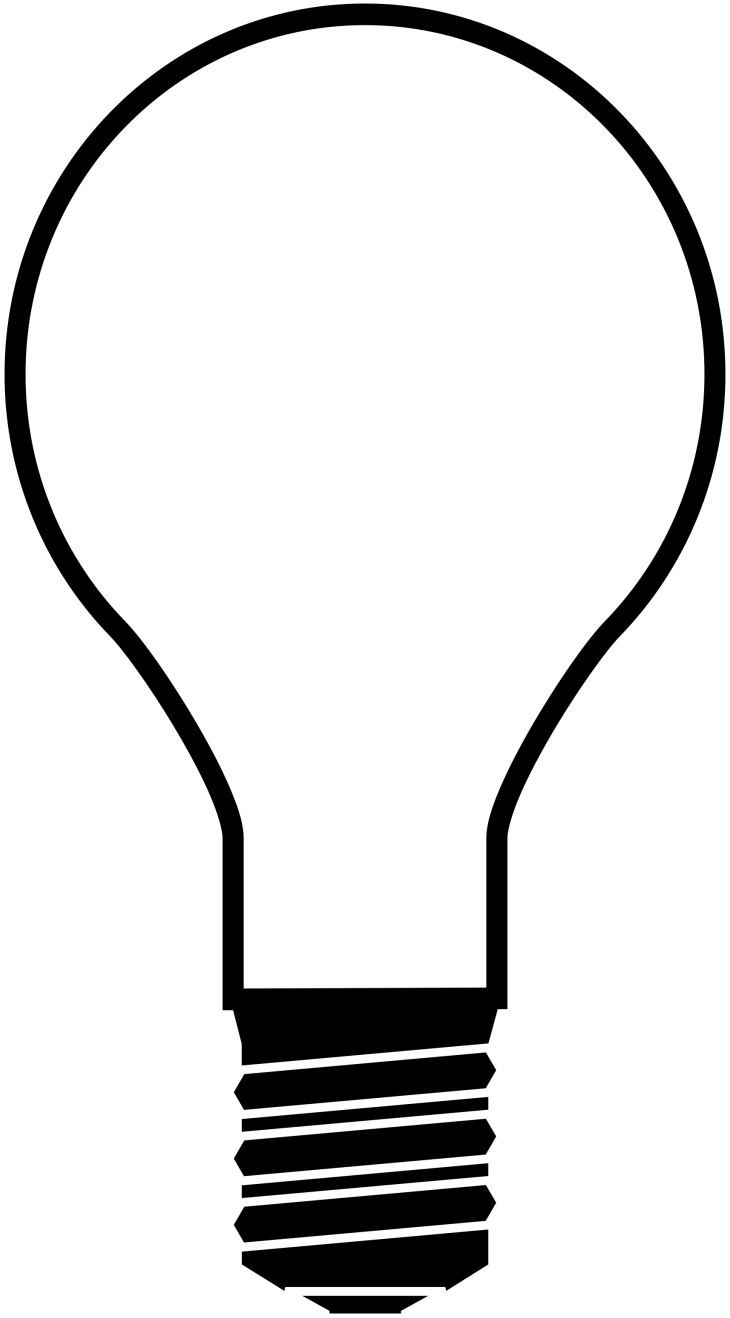

MOVING BEYOND KIDNEYS: LIVERS [Ergin, Sönmez, Ünver *w.p. 2015*]

Similar matching problem (mathematically)

Right Lobe
Segments 5-8

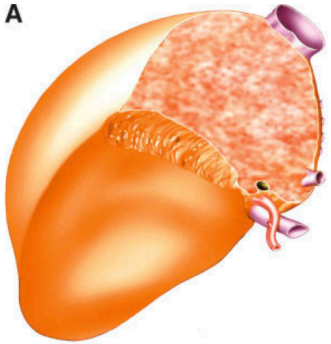

Donor Mortality: 0.5%
Size: 60%
Most risky!

Left + Caudate Lobes
Segments 2-4

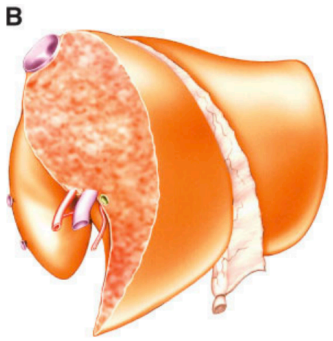

Donor Mortality: 0.1%
Size: 40%
Often too small

Left Lateral Segment
Segments 2-3

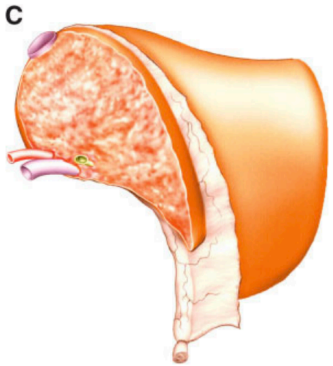

Donor Mortality: Rare
Size: 20%
Only pediatric

[Sönmez 2014]

Right lobe is biggest but riskiest; exchange may reduce right lobe usage and increase transplants

MOVING BEYOND KIDNEYS: MULTI-ORGAN EXCHANGE

[Dickerson Sandholm AAAI-14, JAIR-16]

Chains are great! **[Anderson et al. 2015, Ashlagi et al. 2014, Rees et al. 2009]**

Kidney transplants are "easy" and popular:

- Many altruistic donors

Liver transplants: higher mortality, morbidity:

- (Essentially) no altruistic donors

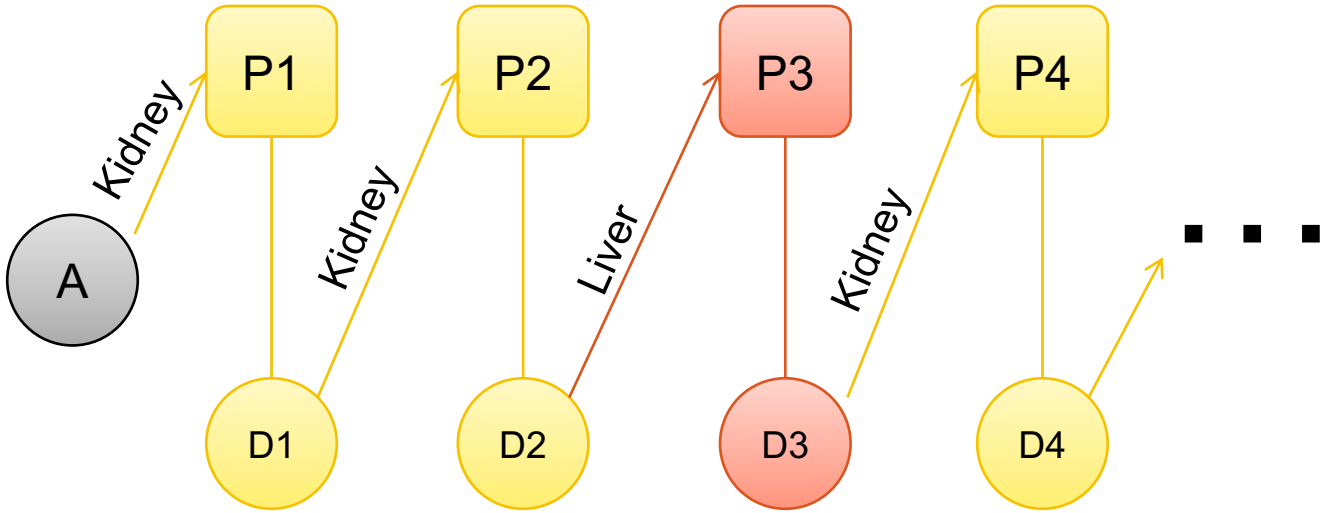

MOVING BEYOND KIDNEYS: LUNGS [Ergin, Sönmez, Ünver *w.p. 2014*]

Fundamentally different matching problem

- **Two** donors needed

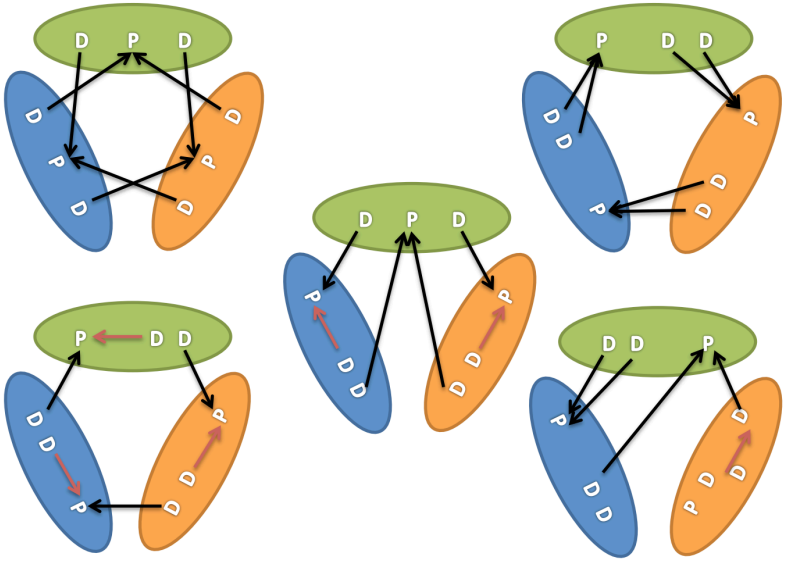

3-way lung exchange configurations

(Compare to the single configuration for a "3-cycle" in kidney exchange.)

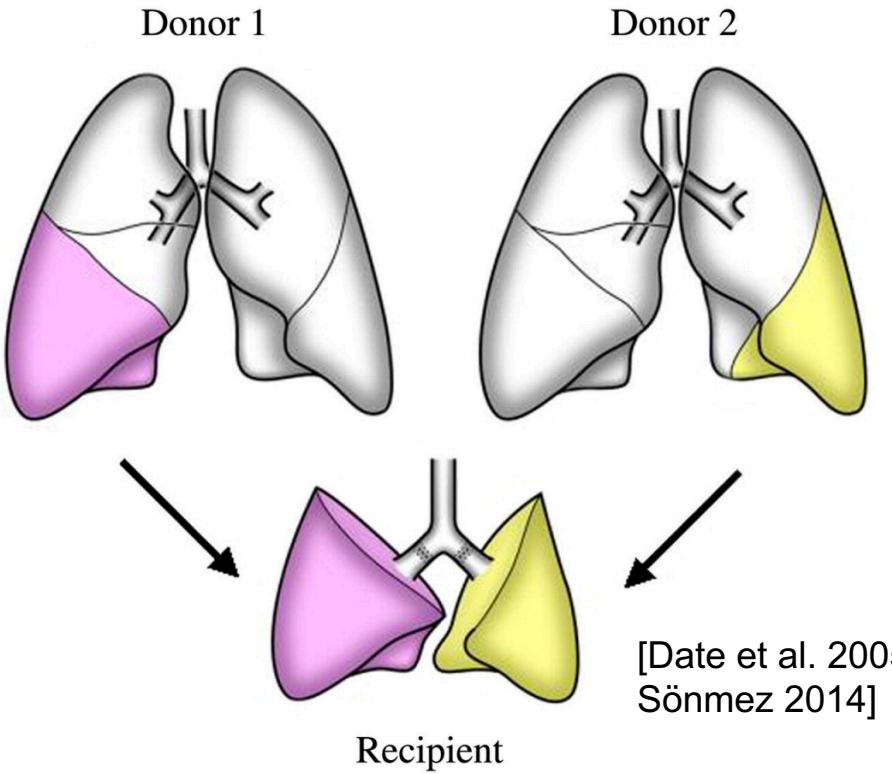

[Date et al. 2005; Sönmez 2014]

NEXT CLASS:

MIDTERM & LIFE QUESTIONS