APPLIED MECHANISM DESIGN FOR SOCIAL GOOD

JOHN P DICKERSON

Lecture #12 – 03/06/2018

CMSC828M
Tuesdays & Thursdays
9:30am – 10:45am

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

THIS CLASS:
ORGAN EXCHANGE

KIDNEY TRANSPLANTATION

- **US waitlist: over 100,000**
  - 36,157 added in 2014
- **4,537 people died while waiting**
- **11,559 people received a kidney from the deceased donor waitlist**
  - (See last class' lecture on deceased donor alloca[illegible]on.)
- **5,283 people received a kidney from a li[illegible] d[illegible]**
  - Some through **kidney excha[illegible]**

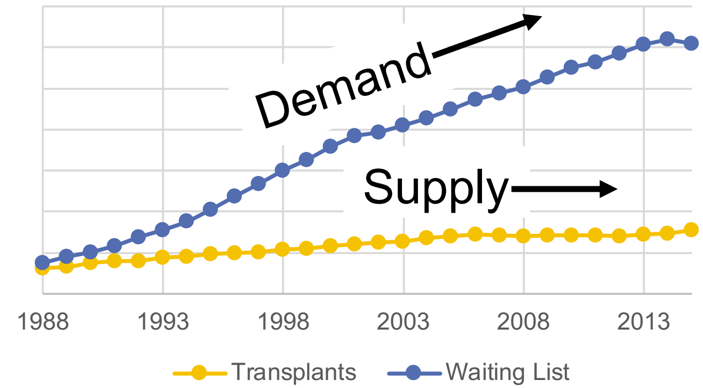

Last time, I promise!

KIDNEY EXCHANGE

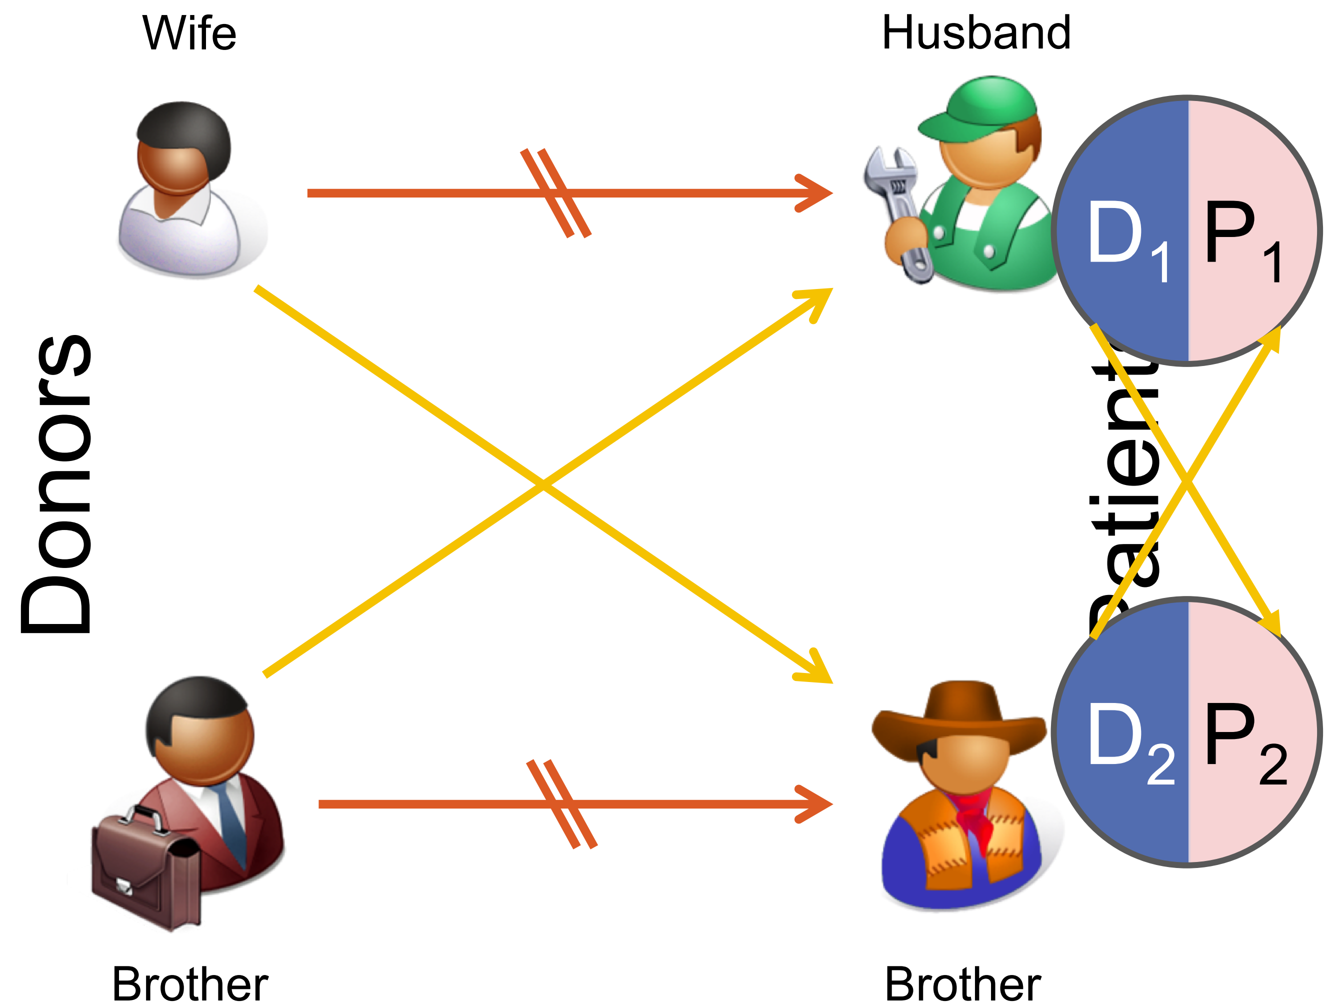

(2- and 3-cycles, all surgeries performed simultaneously)

NON-DIRECTED DONORS & CHAINS

[Rees et al. 2009]

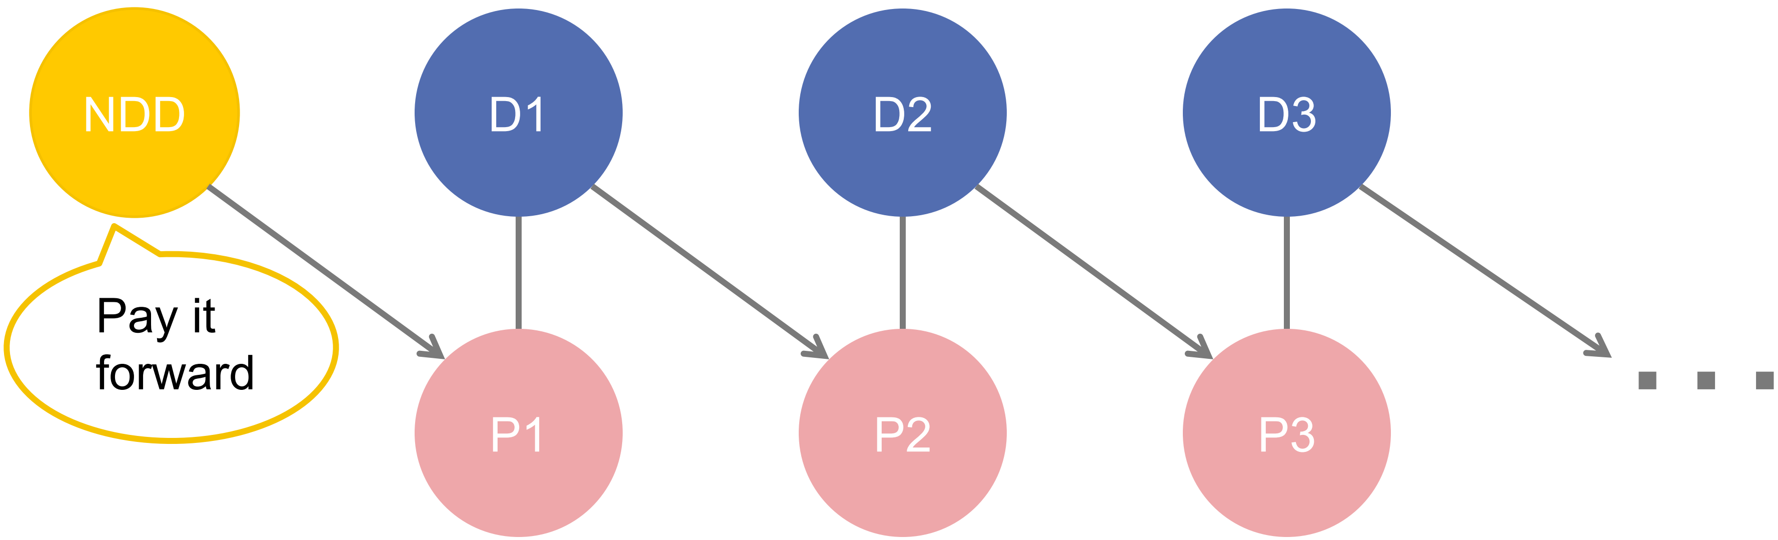

Not executed simultaneously, so no length cap required based on logistic concerns …

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

THE CLEARING PROBLEM

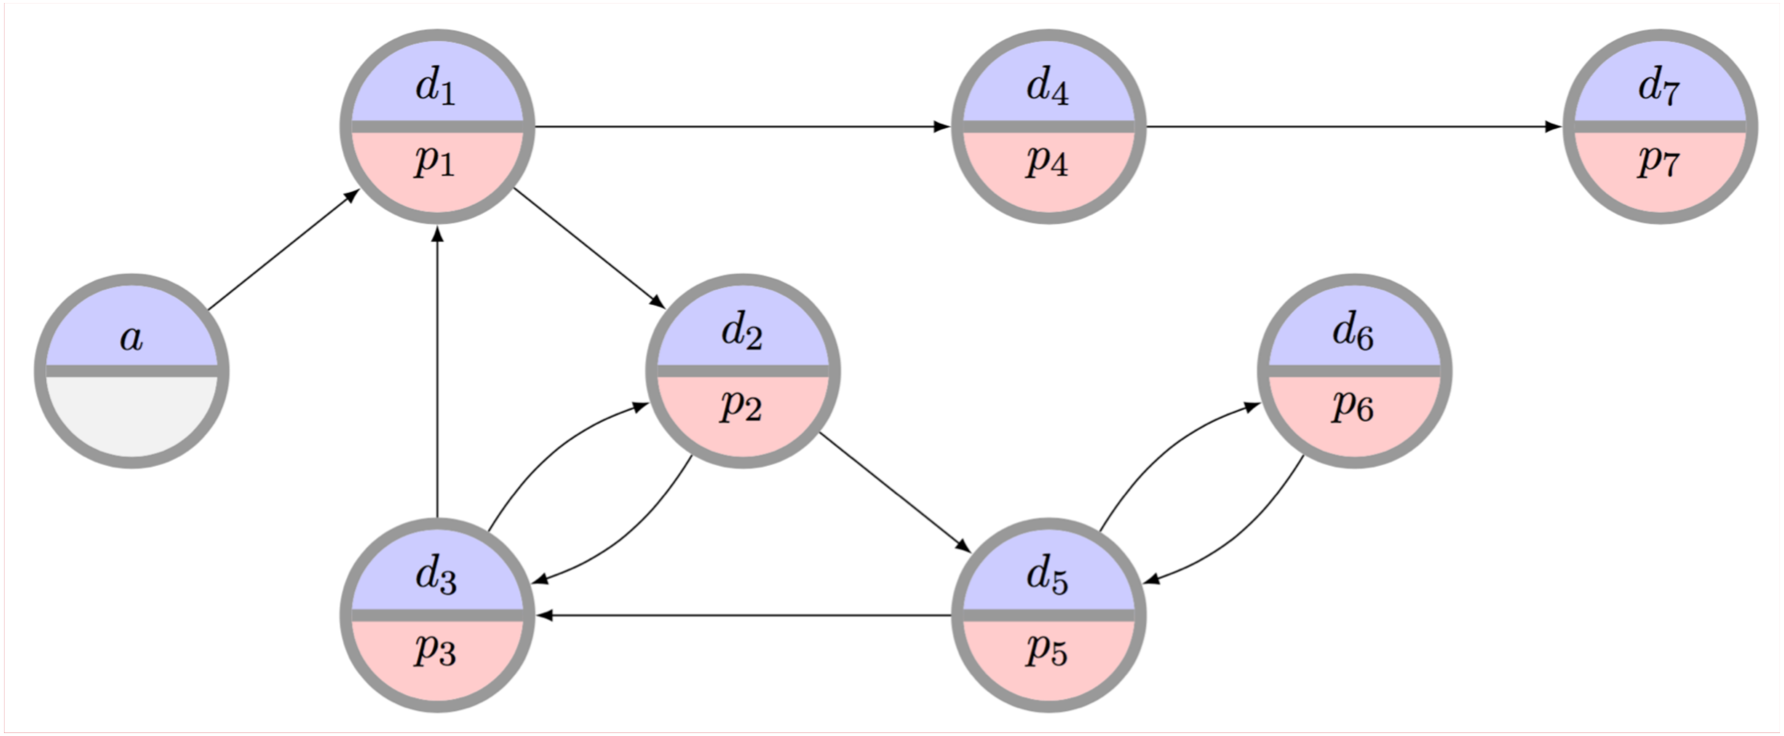

The clearing problem is to find the "best" disjoint set of cycles of length at most *L*, and chains (maybe with a cap *K*)

- Very hard combinatorial optimization problem that we will focus on in the succeeding two lectures.

INDIVIDUAL RATIONALITY (IR)

Will I be better off participating in the mechanism than I would be otherwise?

Long-term IR:

- In the long run, a center will receive at least the same number of matches by participating

Short-term IR:

- At each time period, a center receives at least the same number of matches by participating

STRATEGY PROOFNESS

Do I have any reason to lie to the mechanism?

In any state of the world …

- { time period, past performance, competitors' strategies, current private type, etc }

… a center is not worse off reporting its full private set of pairs and altruists than reporting any other subset

→ No reason to strategize

EFFICIENCY

Does the mechanism result in the absolute best possible solution?

Efficiency:

- Produces a maximum (i.e., max global social welfare) matching given all pairs, regardless of revelation

IR-Efficiency:

- Produces a maximum matching constrained by short-term individual rationality

FIRST: ONLY CYCLES (NO CHAINS)

THE BASIC KIDNEY EXCHANGE GAME [Ashlagi & Roth 2014, and earlier]

Set of *n* transplant centers $T_n = \{t_1 \ldots t_n\}$, each with a set of incompatible pairs V_h

Union of these individual sets is *V*, which induces the underlying compatibility graph

Want: all centers to participate, submit full set of pairs

An allocation *M* is *k*-maximal if there is no allocation *M'* that matches all the vertices in *M* and also more

- Note: *k*-efficient → *k*-maximal, but not vice versa

INDIVIDUALLY RATIONAL?

[Ashlagi & Roth 2014, and earlier]

- **Vertices *a_1*, *a_2* belong to center *a*,**
  b_1, b_2 belong to center *b*
- **Center *a* could match 2 internally**

?????????????????

- **By participating, matches only 1 of its own**
- **Entire exchange matches 3 (otherwise only 2)**

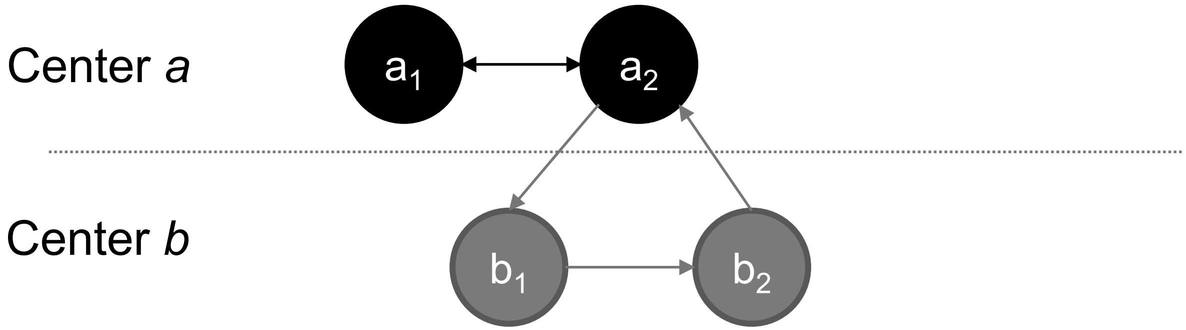

IT CAN GET MUCH WORSE

[Ashlagi & Roth 2014, and earlier]

Theorem: For k>2, there exists *G* s.t. no IR *k*-maximal mechanism matches more than 1/(*k*-1)-fraction of those matched by *k*-efficient allocation

- **Bound is tight**
- **All but one of *a*'s vertices is part of another length *k* exchange (from different agents)**
- ***k*-maximal and IR if *a* matches his *k* vertices (but then nobody else matches, so *k* total)**
- ***k*-efficient to match (k-1)*k**

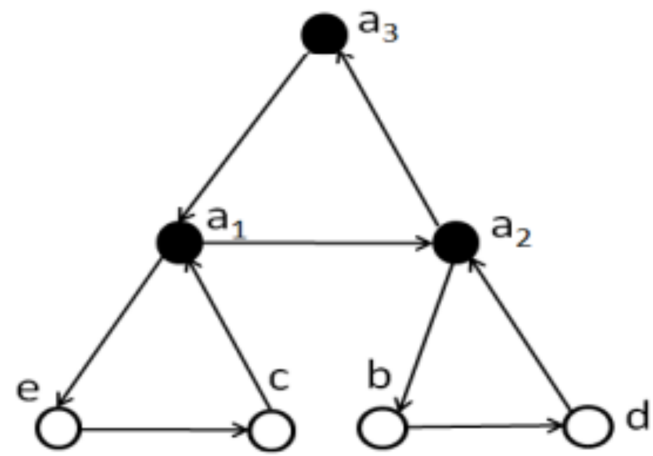

Example: k=3

RESTRICTION #1 [Ashlagi & Roth 2014, and earlier]

Theorem: For all k and all compatibility graphs, there exists an IR k-maximal allocation

Proof sketch: construct *k*-efficient allocation for each specific hospital's pool V_h

Repeatedly search for larger cardinality matching in an entire pool that keeps all already-matched vertices matched (using augmenting matching algorithm from Edmonds)

Once exhausted, done

RESTRICTION #2 [Ashlagi & Roth 2014, and earlier]

Theorem: For k=2, there exists an IR 2-efficient allocation in every compatibility graph

Idea: Every 2-maximal allocation is also 2-efficient

- This is a PTIME problem with, e.g., a standard $O(|V|^3)$ bipartite augmenting paths matching algorithm

By Restriction #1, 2-maximal IR always exists → this 2-efficient IR always exists

RESTRICTION #3 [Ashlagi et al. 2015]

Theorem: No IR mechanism is both maximal and strategyproof (even for k=2)

Suppose mechanism is IR and maximal . . .

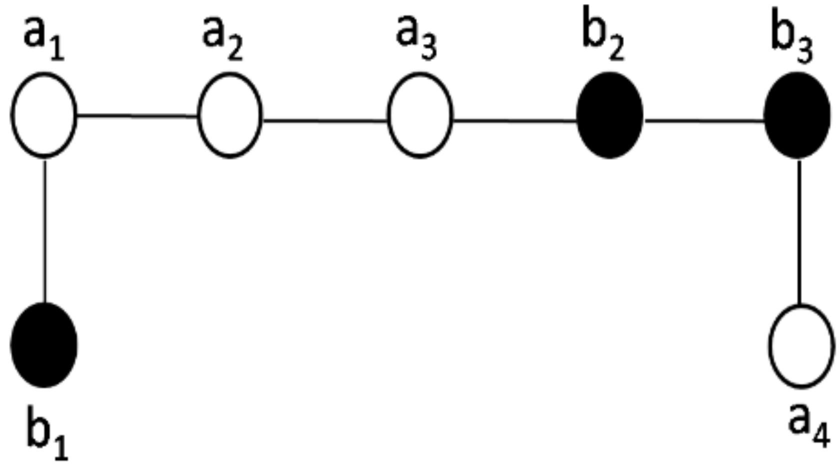

MORE NEGATIVE MECHANISM DESIGN RESULTS [Ashlagi et al. 2015]

Just showed IR + strategyproof → not maximal

No IR + strategyproof mechanism can guarantee more than ½-fraction of efficient allocation

- Idea: same counterexample, note either the # matched for hospital a $\leq$ 3, or # matched for hospital b $\leq$ 2. Proof by cases follows

No IR + strategyproof randomized mechanism can guarantee 7/8-fraction of efficiency

- Idea: same counterexample, bounds on the expected size of matchings for hospitals a, b

HOPELESS ...?

DYNAMIC, CREDIT-BASED MECHANISM [Hajaj et al. AAAI-2015]

Repeated game

Centers are risk neutral, self interested

Transplant centers have (private) sets of pairs:

- Maximum capacity of $2k_i$
- General arrival distribution, mean rate is k_i
- Exist for one time period

Centers reveal subset of their pairs at each time period, can match others internally

CREDITS

Clearinghouse maintains a credit balance c_i for each transplant center over time

High level idea:

- **REDUCE** c_i: center *i* reveals fewer than expected
- **INCREASE** c_i: center *i* reveals more than expected

- **REDUCE** c_i: mechanism tiebreaks in center *i*'s favor
- **INCREASE** c_i: mechanism tiebreaks against center *I*

Also remove centers who misbehave "too much."

Credits now → matches in the future

THE DYNAMIC MECHANISM

1. **Initial credit update**
   - Centers reveal pairs
   - Mechanism updates credits according to k_i
2. **Compute maximum global matching**
   - Gives the utility U_g of *a* max matching
3. **Selection of a final matching**
   - Constrained to those matchings of utility U_g
   - Take c_i into account to (dis)favor utility given by matching to a specific center i
   - Update c_i based on this round's (dis)favoring
4. **Removal phase if center is negative for "too long"**

THEORETICAL GUARANTEES

Theorem: No mechanism that supports cycles and chains can be both long-term IR and efficient

Theorem: Under reasonable assumptions, the prior mechanism is both long-term IR and efficient

LOTS OF OPEN PROBLEMS HERE

Dynamic mechanisms are more realistic, but …

- Vertices disappear after one time period
- All hospitals the same size
- No weights on edges
- No uncertainty on edges or vertices
- Upper bound on number of vertices per hospital
- Distribution might change over time
- ...

Project?

WHAT DO EFFICIENT MATCHINGS EVEN LOOK LIKE …?

Next class: given a specific graph, what is the "optimal matching"

This class: given a family of graphs, what do "optimal matchings" tend to look like?

Use a stylized random graph model, like [Saidman et al. 2006]:

- Patient and donor are drawn with blood types randomly selected from PDF of blood types (roughly mimics US makeup), randomized "high" or "low" CPRA
- Edge exists between pairs if candidate and donor are ABO-compatible and tissue type compatible (random roll weighted by CPRA)

RANDOM GRAPH PRIMER

Canonical Erdős-Rényi random graph *G*(*m*,*p*) has m vertices and an (undirected) edge between two vertices with probability *p*

- Let *Q* be the property of "there exists a perfect matching" in this graph

The convergence rate to 1 (i.e., "there is almost certainly a near perfect matching in this graph) is exponential in *p*

- $\Pr(G(m,m,p)$ satisfies $Q) = 1 - o(2^{-mp})$
- At least as strong with non-bipartite random graphs

Early random graph results in kidney exchange are for "in the large" random graphs that (allegedly) mimic the real compatibility graphs

- All models are wrong, but some are useful?

A STYLIZED ERDŐS-RÉNYI-STYLE MODEL OF KIDNEY EXCHANGE

In these random (ABO- & PRA-) graphs:

- # of O-{A, B, AB} pairs > {A, B, AB}-O pairs
- # of {A, B}-AB pairs > AB-{A, B} pairs
- Constant difference between # A-B and # B-A

Idea #1: O-candidates are hard to self-match

Idea #2: {A, B}-candidates are hard to self-match

Idea #3: "symmetry" between A-B and B-A (equally hard to self-match, give or take)

EFFICIENT MATCHING IN DENSE GRAPHS WITH ONLY CYCLES

Under some other assumptions about PRA …

Almost every large random (ABO- & PRA-) graph has an efficient allocation that requires exchanges of size at most 3 with the following:

- X-X pairs are matched in 2- or 3-way exchanges with other X-X pairs (so-called “self-demand”)
- B-A pairs are 2-matched with A-B pairs
- The leftovers of {A-B or B-A} are 3-matched with “good” {O-A, O-B} pairs and {O-B, O-A pairs}
- 3-matches with {AB-O, O-A, A-AB}
- All the remaining 2-matched as {O-X, X-O}

VISUALLY …

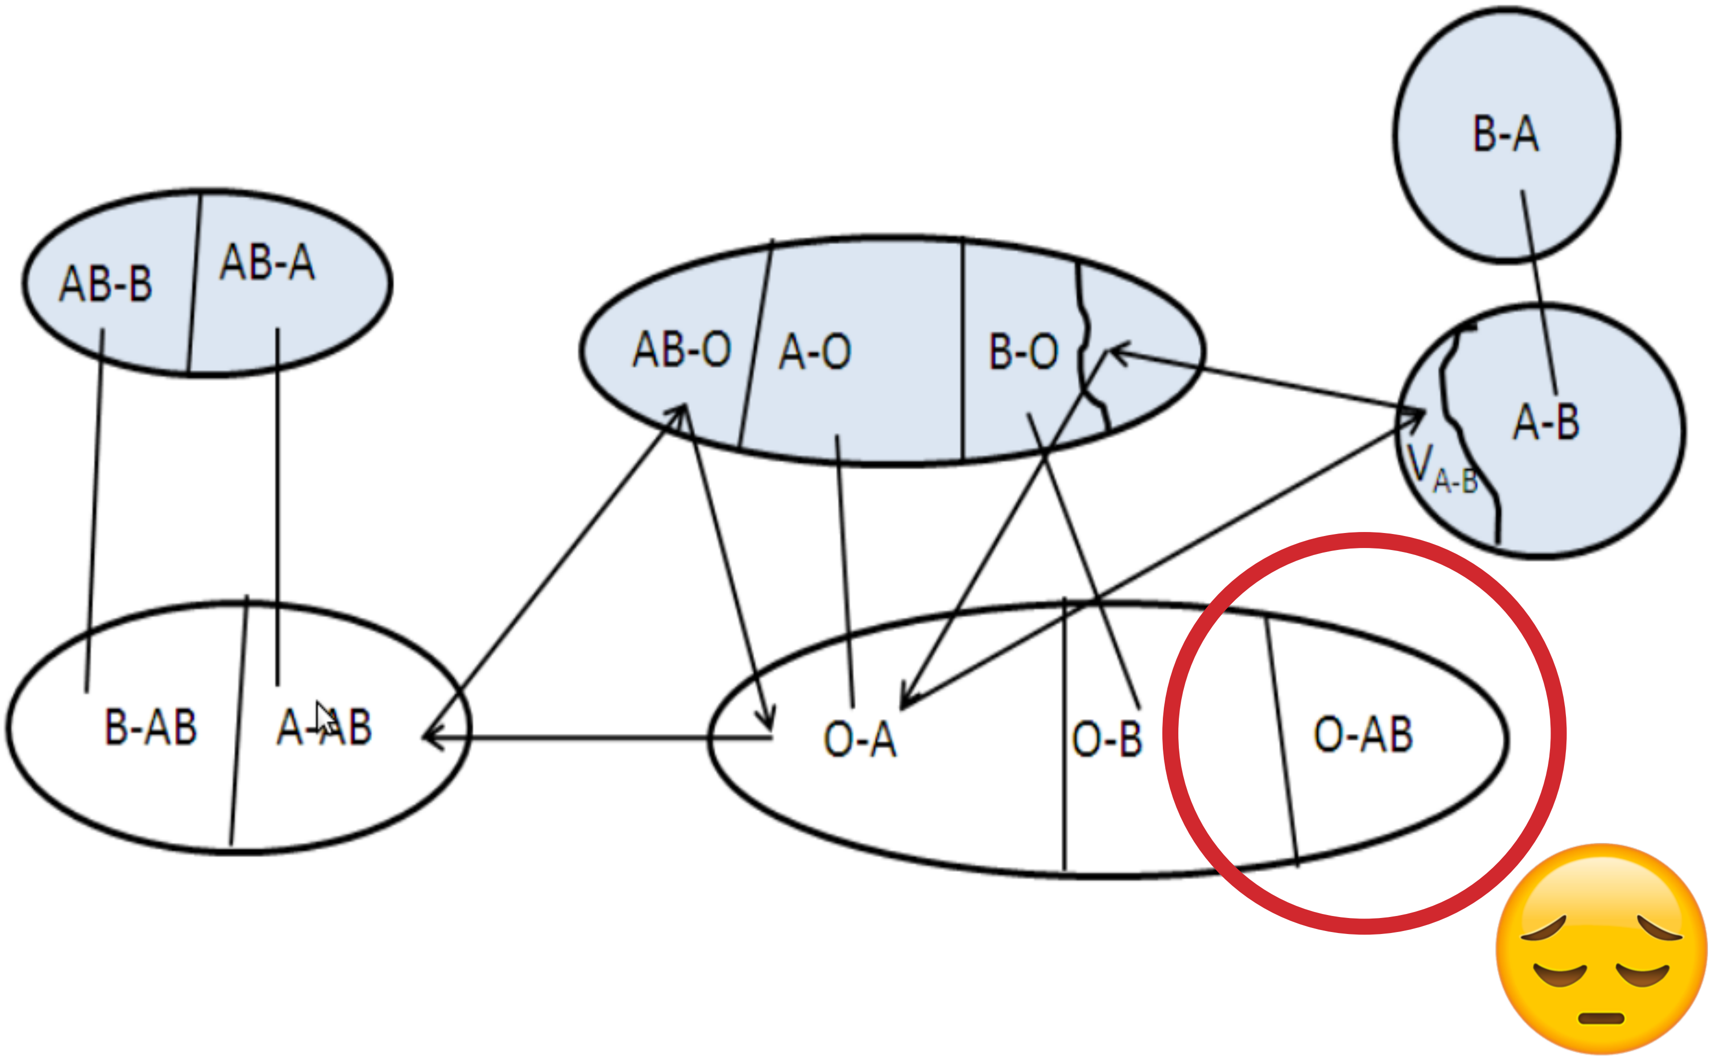

PRICE OF FAIRNESS

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

[Dickerson et al. AAMAS-14]

- **Clearing problem**: find a matching M^* that maximizes utility function

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

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

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

$V_{\{L,H\}}$ **: lowly-, highly-sensitized vertices**

λ : fraction of pool that is lowly-sensitized

$p_{\{L,H\}}$ **: prob. ABO-compatible is tissue-type incompatible**

$p = \lambda p_L + (1-\lambda)p_H$ **: average level of sensitization**

"Most stringent" fairness rule:

$$u_{H \succ L}(M) = \begin{cases} u(M) & \text{if } |M_H| = \max_{M' \in \mathcal{M}} |M'_H| \\ 0 & \text{otherwise} \end{cases}$$

Theorem

Assume p < 2/5, λ ≥ 1-p, and "reasonable" distribution of blood types.
Then, almost surely as $n \to \infty$,

$$\mathrm{POF}(\mathcal{M}, u_{H \succ L}) \leq \frac{2}{33}.$$

(And this is achieved using cycles of length at most 3.)

IN THEORY, THE PRICE OF FAIRNESS IS LOW

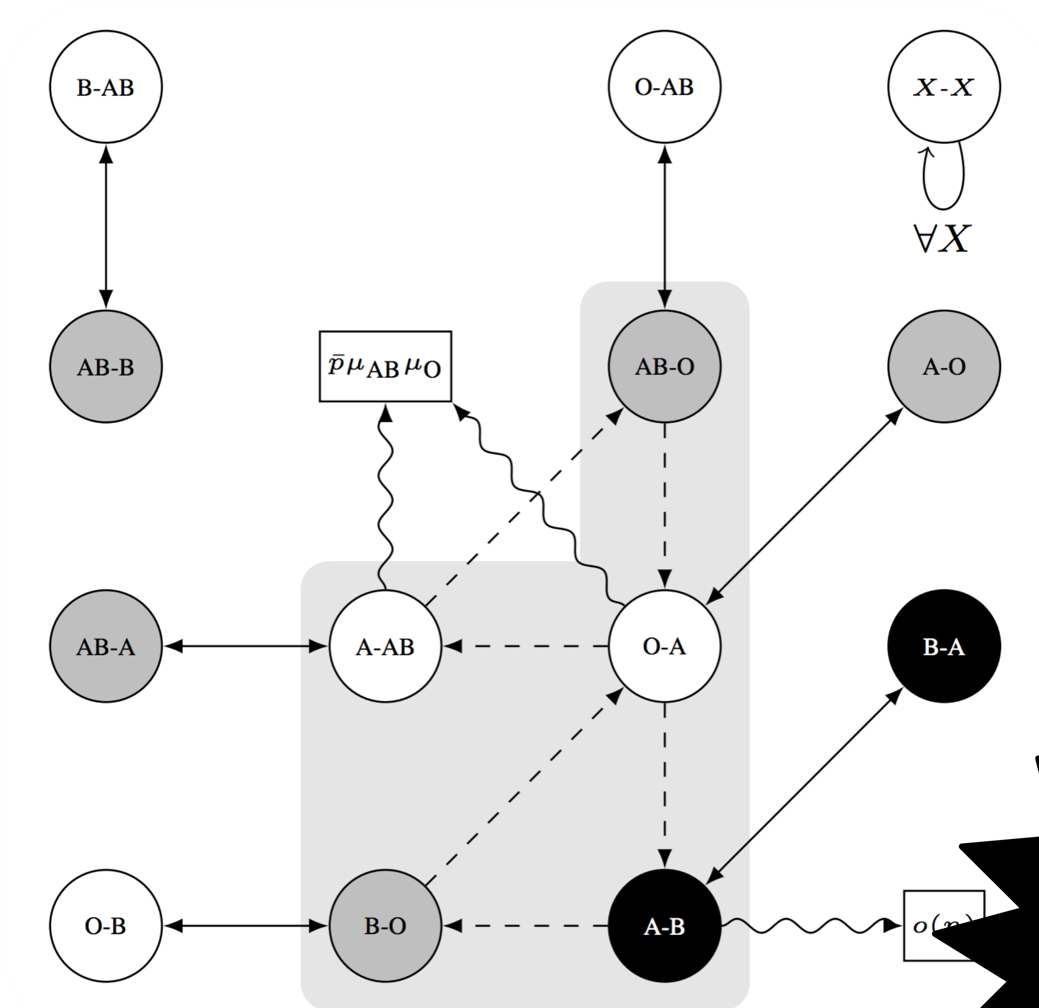

More on 10/6!

PROBLEMS WITH THIS TYPE OF MODEL

Dense model [Saidman et al. 2006, etc.]

- Constant probability of edge existing
- Less useful in practice [Ashlagi et al. 2012+, Dickerson et al. 2014+]

Better? Sparse model [Ashlagi et al. 2012]

- 1-λ fraction is *highly-sensitized* ($p_H = c/n$)
- λ fraction is *lowly-sensitized* ($p_L > 0$, constant)

But still:

- Random graph models tend to be "in the large", no weights, no uncertainty, fairly homogeneous … so not perfect!

A TASTE OF THE SPARSE MODEL ...

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

Similar matching problem (mathematically)

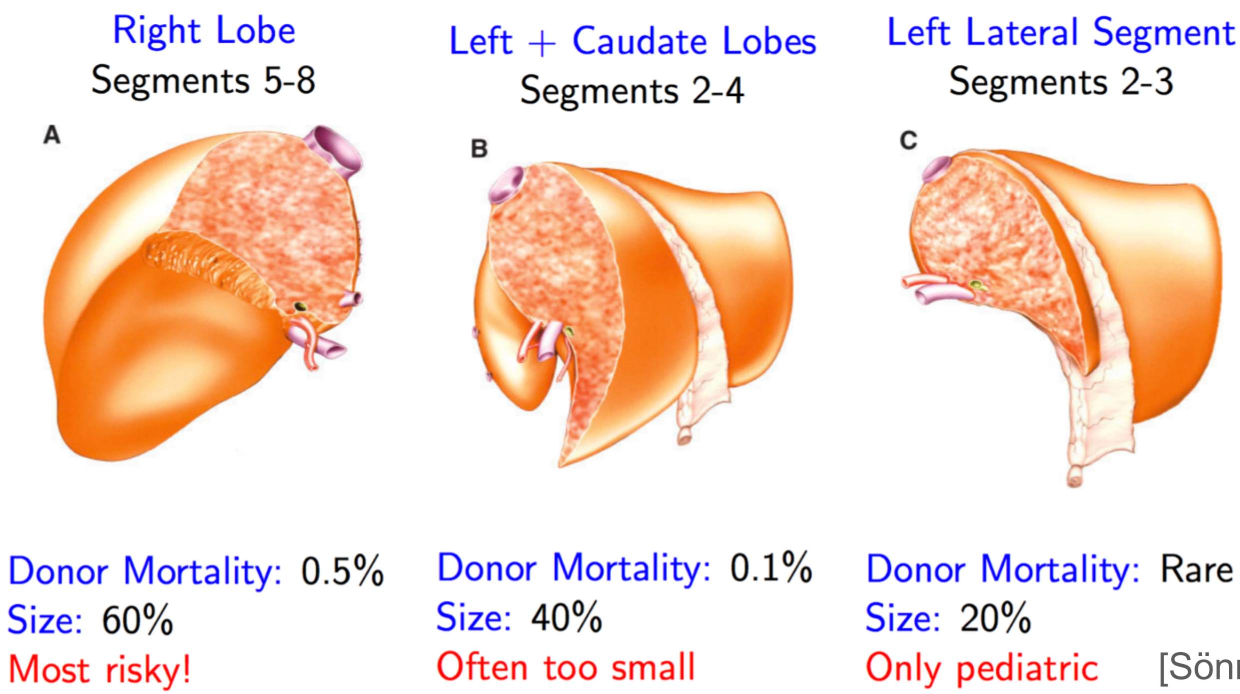

[Sönmez 2014]

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

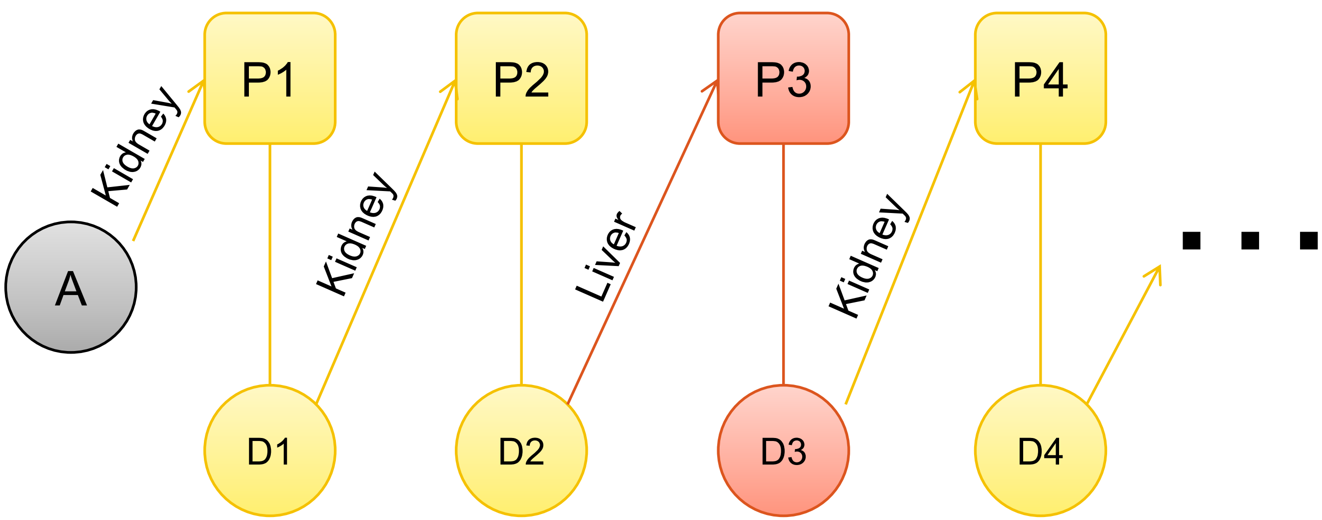

SPARSE GRAPH, MANY ALTRUISTS

n_K kidney pairs in graph D_K; $n_L = \gamma n_K$ liver pairs in graph D_L

Number of altruists t(n_K)

Constant $p_{K \rightarrow L} > 0$ of kidney donor willing to give liver

Constant cycle cap z

Theorem

Assume $t(n_K) = \beta n_K$ for some constant $\beta > 0$. Then, with probability 1 as $n_K \rightarrow \infty$,

Any efficient matching on $D = \text{join}(D_K, D_L)$ matches $\Omega(n_K)$ more pairs than the aggregate of efficient matchings on D_K and D_L.

Building on [Ashlagi et al. 2012]

INTUITION

Find a linear number of "good cycles" in D_L that are length > z

- Good cycles = isolated path in highly-sensitized portion of pool and exactly one node in low portion

Extend chains from D_K into the isolated paths (aka can't be matched otherwise) in D_L, of which there are linearly many

- Have to worry about $p_{K \to L}$, and compatibility between vertices

Show that a subset of the dotted edges below results in a linear-in-number-of-altruists max matching

- → linear number of D_K chains extended into D_L
- → linear number of previously unmatched D_L vertices matched

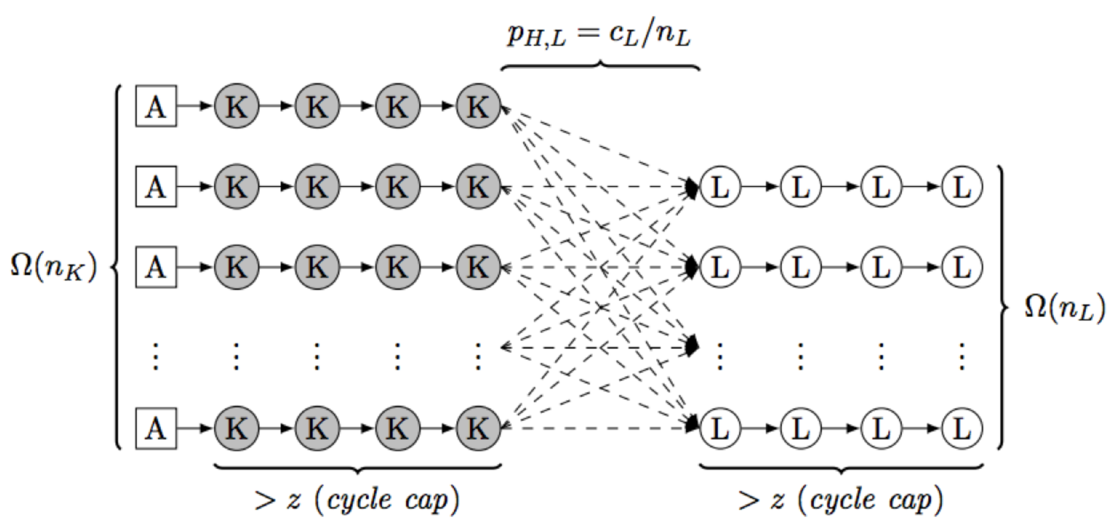

SPARSE GRAPH, FEW ALTRUISTS

n_K kidney pairs in graph D_K; $n_L = \gamma n_K$ liver pairs in graph D_L

Number of altruists t – no longer depends on n_K!

λ is frac. lowly-sensitized

Constant cycle cap z

Theorem

Assume constant t. Then there exists λ' > 0 s.t. for all λ < λ'

Any efficient matching on D = join(D_K,D_L) matches $\Omega(n_K)$ more pairs than the aggregate of efficient matchings on D_K and D_L.

With constant positive probability. *Building on [Ashlagi et al. 2012]*

INTUITION

For large enough λ (i.e., lots of sensitized patients), there exist pairs in D_K that can't be matched in short cycles, thus only in chains

- Same deal with D_L, except there are no chains

Connect a long chain (+altruist) in D_K into an unmatchable long chain in D_L, such that a linear number of D_L pairs are now matched

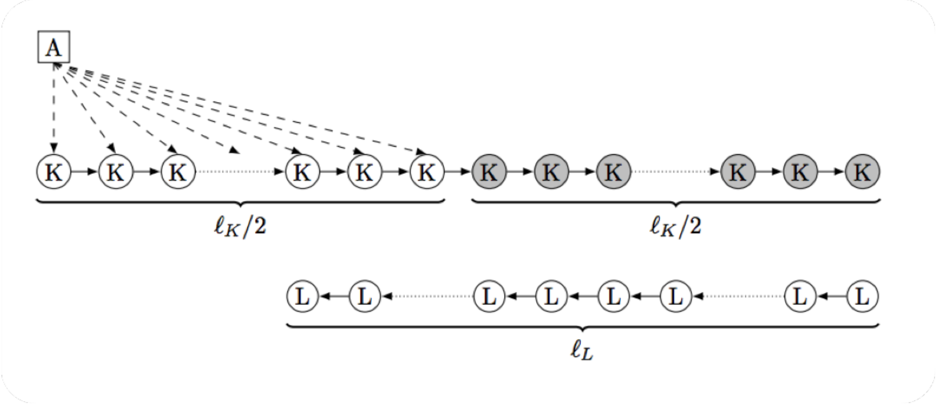

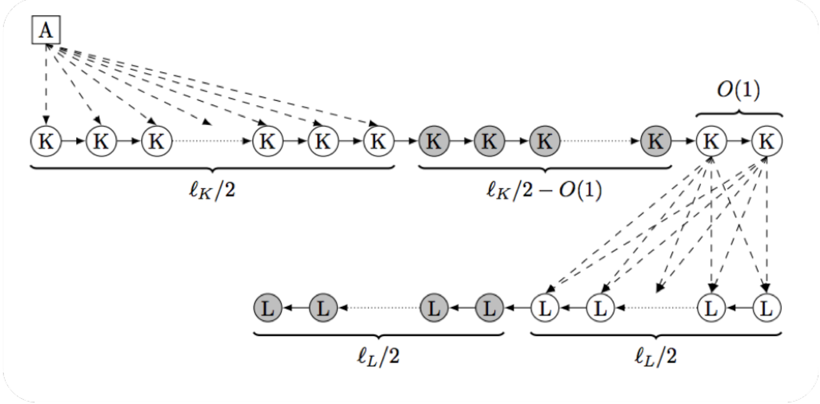

MOVING BEYOND KIDNEYS: LUNGS [Ergin, Sönmez, Ünver *w.p. 2014*]

Fundamentally different matching problem

- **Two** donors needed

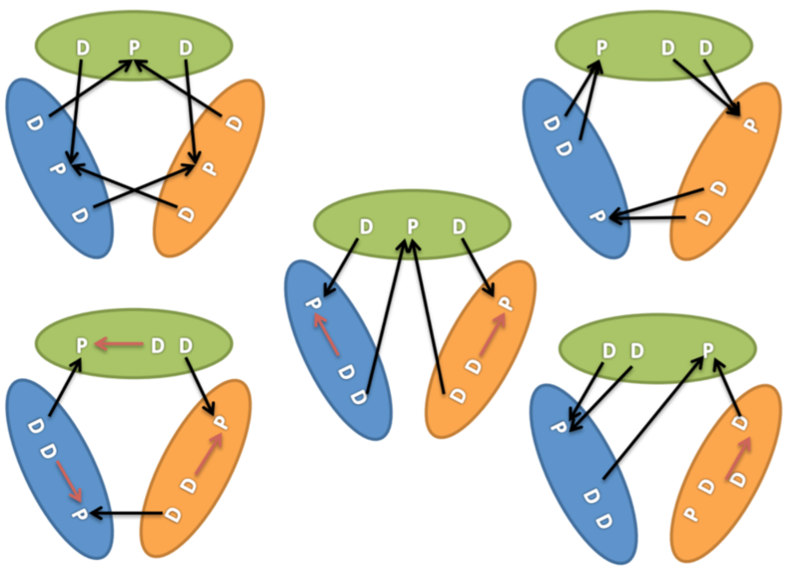

3-way lung exchange configurations

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

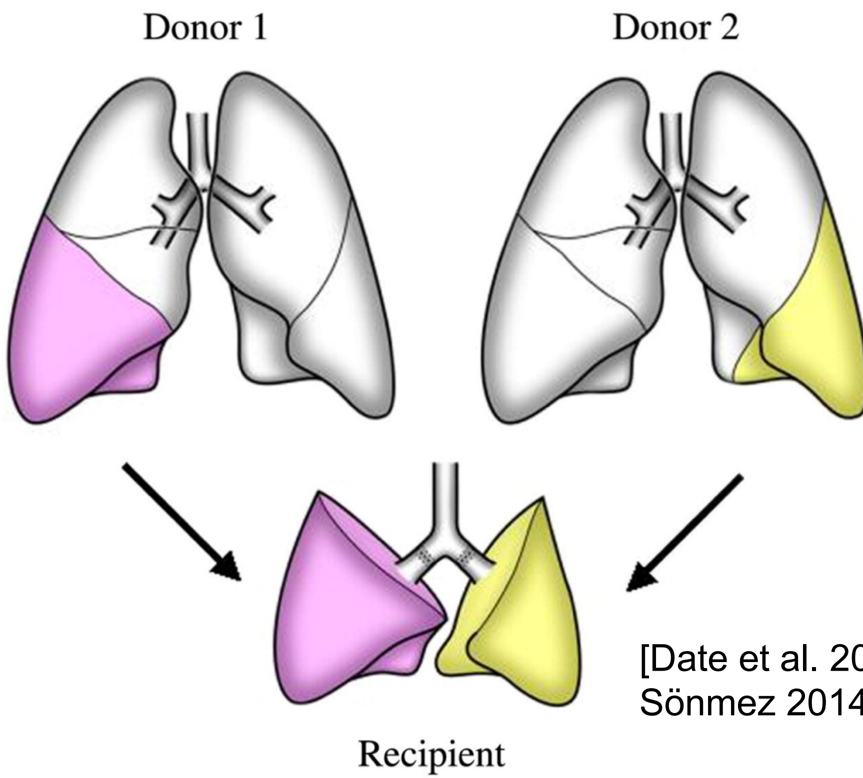

[Date et al. 2005; Sönmez 2014]

OTHER RECENT & ONGOING RESEARCH IN THIS SPACE

Dynamic matching theory with a kidney exchange flavor:

- Akbarpour et al., "Thickness and Information in Dynamic Matching Markets"
- Anderson et al., "A dynamic model of barter exchange"
- Ashlagi et al., "On matching and thickness in heterogeneous dynamic markets"
- Das et al., "Competing dynamic matching markets"

Mechanism design:

- Blum et al. "Opting in to optimal matchings"

Not "in the large" random graph models:

- Ding et al., "A non-asymptotic approach to analyzing kidney exchange graphs

NEXT CLASS:
OPTIMAL BATCH CLEARING OF ORGAN EXCHANGES