Flexible Multi-unit Exchange:
Theory and An Application to Blood Allocation with Replacement Donors

Xiang Han
Shanghai UFE

Onur Kesten
University of Sydney

M. Utku Ünver
Boston College

University of Tokyo
Market Design Center
June 18, 2026

Transfusion

- **Blood components** — red blood cells, platelets, plasma, and **whole blood** itself — are the most used human body components in medical treatments.
  - They are mostly used to replace or supplement missing components from blood.
  - Whole blood is rarely used in modern transfusion medicine.
- Transfusion is one of the most practiced medical procedures in hospitals (most common inpatient procedure in U.S. in 2011).
- There are severe blood shortages in many developing countries and seasonal shortages in developed countries.
  - Annual shortage in low- and middle-income countries (LMICs): ∼102 million units (worldwide donations were only 118.5 million in 2018).
  - Millions of deaths each year are associated with lack of blood.

Donation

- World Health Organization (WHO) strongly recommends all blood is supplied by volunteer non-remunerated donors (VNRDs) due to safety and ethics concerns.
- Donating blood components is substantially less costly than any other organ or tissue donation procedure for a donor.
  - A donor can donate red blood cells or whole blood once in every 56 days, while donation of platelets and plasma can be made more often through a procedure called apheresis.
- However,
  - VNRDs are insufficient to meet the demand for blood in many regions of the world.
  - It is costly to screen VNRDs for safety concerns and quality for developing countries, especially in sub-Saharan Africa.

Donation

- Latest WHO reports (WHO'25):
  - only 79 countries (38 high-, 33 middle-, 8 low-income) collect more than 90% of blood from VNRDs; and
  - at least 55 countries rely heavily on replacement donors (RDs) — more than 50% of national supply (e.g., Pakistan, Nigeria, Bangladesh, Mexico; pop. $>$ 1.5 bn).
- WHO data appear to underestimate RD use (e.g., Kenya, Venezuela reported as 100% VNRD; local studies show most blood is from RDs).

Replacement Donor Programs

- In a replacement donor (RD) program, a patient receives the blood she needs by bringing forward a number of replacement donors — usually family members or loved ones of the patient.
- Blood banks use different exchange rates that determine how many units must be supplied for each unit received, depending on:
  - shortages at the blood bank,
  - patients' medical urgency, and
  - other demographics of the patient.
- RD programs are observed in all continents, and especially common in Africa, Latin America, and Central Asia (Allain & Sibinga'16).

RD Systems within the Broader Blood Market

- Demand: non-emergency use dominates.
  - Africa: only 8.5% of units went to emergency departments in 2020.
  - RD programs target non-emergency patients; bank inventory covers emergencies.
- Supply: the natural hybrid model.
  - VNRDs — centralized; RD programs — decentralized, hospital-based.
  - Two channels organized independently.

Implication: we can isolate the RD submarket and treat VNR inventory as exogenous.

Replacement Donation Is Not Going Away

- Many experts and local authorities argue that RDs are as safe and ethical as VNRDs, and the only ethically acceptable and economically affordable option for mitigating shortages in many countries.
- Economic, institutional, and cultural factors have consistently delayed the development of VNRD systems and made replacement donation a durable and often dominant source of blood in many countries.
- Despite 50+ years of WHO calls for robust VNRD systems, the global share of RD blood in 2018 was almost unchanged from 2002, while the absolute amount increased by 25%.

Replacement Donor Programs

Table 1: Whole Blood Donations Reported by Countries to WHO (millions)

	Total Donations	VNR Donations	% of VNR Donations	FR Donations	% of FR Donations	Paid Donations	% of Paid Donations	Other (unknown)	% of Other
1998-1999 (n=175)	75.13	55.9	74.404	13.5	17.969	5.73	7.627		
2001-2002 (n=178)	76.27	61.9	81.159	12.1	15.865	2.27	2.976		
2004-2005 (n=154)	64.824	51.52	79.477	11.03	17.015	2.274	3.508		
2008 (n=159)	65.51	54.34	82.949						
2013 (n=178)	88.2	73.471	83.3	14.465	16.4	0.265	0.3		
2018 (n=171)	95.228	78.8	82.749	15.1	15.857	0.16	0.168	1.168	1.227

- Global share of blood supply from replacement donation in 2018 ($\sim$15.9%) is almost unchanged from 2001–2002, but the absolute amount has grown substantially.

Where RDs Dominate the Blood Supply

- For 53 countries with detailed data that rely heavily on replacement donation (more than 50% of blood supply from RDs) in 2018:
  - 78.05% of total blood supply comes from RDs, slightly higher than 2011 (77.34%),
  - compared with 2011, total replacement donations increased by 48.52%, and
  - total population of these countries is over 1.5 billion.
- Implication: in many LMICs, designing better RD programs has first-order welfare consequences — not a niche intervention.

Shortcomings of Existing RD Programs

Existing programs suffer from two major shortcomings:

- They lack optimization of allocation.
  - De facto first-come, first-served (FCFS): patients only exchange with the blood bank, not with each other.
  - Can be highly inefficient when bank inventory is thin.
  - No donor screening in general, depending on the needs of the blood bank.
- They use exogenously fixed exchange rates, creating efficiency, fairness, and ethics problems.
  - Some patients cannot recruit the required number of donors.
  - Gives rise to coercion and black markets in which third parties are paid to serve as RDs.
  - Rules are bent arbitrarily: some patients are forgiven their *debts*, others are not (e.g., Tucumán, Argentina). Trust and transparency are eroded.

Our Approach

- We model and design multi-unit exchange under flexible exchange rates — with the leading application being blood allocation with VNRDs and RDs.
- We introduce a new policy lever: feasible schedule menus (FSMs) — they formalize potentially endogenous exchange rates between units a patient receives and supplies.
- We introduce a broad class of weighted utilitarian mechanisms (WUMs).
  - Encompass priority mechanisms and maximal mechanisms.
- We discipline FSMs by a discrete-convexity condition — $L^\natural$-convexity (L-natural-convexity, Murota'98) — and two other natural conditions.

Main Results and Contributions

- Under suitable FSMs, WUMs satisfy:
  - Pareto efficiency and an approximate equal treatment of equals;
  - incentive compatibility with respect to donor revelation under separable utility functions;

  and priority mechanisms (incl. maximal) are fully incentive compatible (donors and utility functions).
- FCFS can be Pareto-dominated by an IC mechanism — under the two most common exchange-rate policies.
- Simulations: large gains over FCFS at low inventory (relevant for thin LMIC markets).
- Practical design: 4-parameter menus, hybrid implementation with VNRDs, fairness and blood-quality designs.
- Results extend to general multi-unit exchange with compatibility-based preferences: time banks, shift exchange.

Some Related Literature

- Incentives & altruism in blood donations: Lacetera, Macis & Slonim'12,'13; Slonim, Wang & Garbarino'14
- Replacement donors: Allain & Sibinga'16; Sun, Lu & Jin'16
- COVID-19 blood-plasma donation, pay-forward/-backward: Kominers, Pathak, Sönmez & Ünver'20
- Living-donor kidney exchange, weighted/maximal/priority: Roth, Sönmez & Ünver'05,'07; Andersson & Kratz'19; Kratz'24;
  dual-organ exchange: Ergin, Sönmez & Ünver'17
- Compatibility-based preferences: Bogomolnaia & Moulin'04; Nicolò, Yadav & Sen'19; Aziz'20; Echenique, Miralles & Zhang'20; Garg, Tröbst & Vazirani'20; Andersson, Cseh, Ehlers & Erlanson'21; Manjunath & Westkamp'21
- Other multi-unit exchange: Moulin'95; Konishi, Quint & Wako'01; Biró, Klijn & Pápai'22; Feng, Klaus & Klijn'24
- Beyond one-for-one: Hylland & Zeckhauser'79; Budish'11; Agarwal et al.'19; Echenique, Goel & Lee'24
- Discrete convex analysis ($M^\natural$/$L^\natural$): Murota'98; Fujishige & Yang'03,'06; Kojima, Tamura & Yokoo'18; Hatfield et al.'19; Hafalir, Kojima & Yenmez'25
- Dominating SP via SP: Erdil'14; Alva & Manjunath'21
- Minimalist market design: Sönmez'24

Blood Transfusion

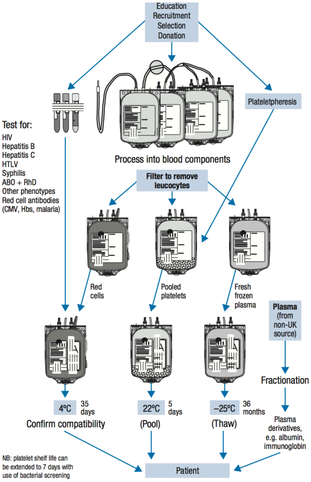

- One donor can donate on average
  - ≈500 ml of whole blood (1 unit of red blood cells – once in every 56 days), or
  - up to 2-3 units of platelets through apheresis, once in a week up to 24 times a year, or
  - up to 4 units of plasma through apheresis, twice a week, there can be a limit depending on the purpose of plasma usage.
- Components are processed and then packaged in standard packs.
- A patient typically needs multiple units of compatible blood component for one time use or multiple uses: Idiosyncratic discrete units in some bounded interval are feasible.
- More blood is, in expectation, better for a patient's prognosis (within the interval).
- Different components typically have different (*ABO* and *RhD*) blood type compatibility requirements.

Model

- I: set of patients in need of a given blood component
- $\mathcal{B}$: set of relevant blood types
  - e.g., $\mathcal{B} = \{O+, O-, A+, A-, B+, B-, AB+, AB-\}$
- Each patient $i \in I$
  - has blood type $\beta_i \in \mathcal{B}$,
  - is compatible with blood types in $\mathcal{C}(\beta_i) \subseteq \mathcal{B}$,
  - needs a maximum of $\overline{n}_i$ units of compatible blood,
  - has a (possibly empty) RD set D_i, where each $d \in D_i$ has blood type $\beta_d \in \mathcal{B}$ and can donate 1 unit.
- b: the blood bank, with an inventory of v_X units of type X blood for each $X \in \mathcal{B}$
- δ: upper bound on the number of donors a patient can bring forward.

(Extends one-for-one frameworks of Andersson, Cseh, Ehlers & Erlanson'21 for time banks and Manjunath & Westkamp'21 for shift exchange)

Schedules

- The schedule of patient i is a pair of nonnegative integers (r, s):
  - r: amount of compatible blood received
  - s: amount of blood supplied (through her donors)
  - All possible schedules for i: $\{0, 1, \ldots, \overline{n}_i\} \times \{0, 1, \ldots, \delta\}$.

($r = s$ in Andersson, Cseh, Ehlers & Erlanson'21 and Manjunath & Westkamp'21; we allow flexible exchange rates.)

Preferences and Utilities

- Patient i has separable preferences over schedules, represented by

$$u_i(r, s) = \rho_i(r) - \sigma_i(s).$$

  - blood valuation function $\rho_i : \{0, 1, \ldots, \overline{n}_i\} \to \mathbb{R}$ is strictly increasing and concave (decreasing marginal utility);
  - cost function for supply $\sigma_i : \{0, 1, \ldots, \delta\} \to \mathbb{R}$ is strictly increasing and convex (increasing marginal cost);
  - for any $0 \le r < \overline{n}_i$ and $0 \le s < \delta$,

  $u_i(r{+}1, s{+}1) > u_i(r, s)$ i.e., marginal rate of substitution (MRS

- A quasi-linear special case: $u_i(r, s) = \rho_i(r) - \theta_i \cdot s$, with $\theta_i > 0$.
- The set of possible utility functions is finite.

Preferences and Utilities

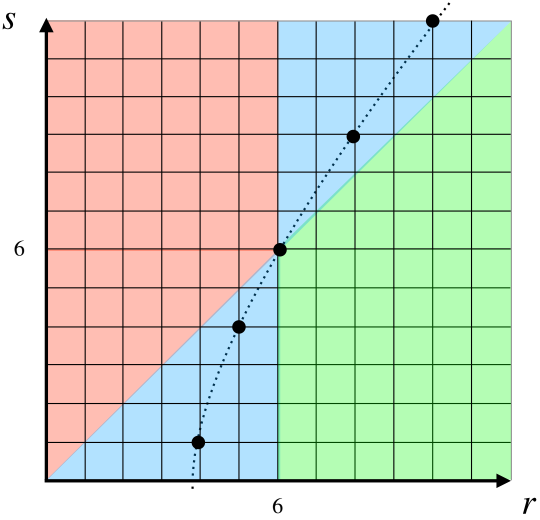

Blood Bank Objectives

- The blood bank cares about leftover blood of each type in its inventory $(r_X)_{X\in\mathcal{B}}$, and possibly the total transfusions to patients, r_t. The transfusion and inventory objective is

$$o\big(w,(r_X)_{X\in\mathcal{B}},r_t\big) \;=\; \sum_{X\in\mathcal{B}} w_X\cdot r_X \;+\; w_t\cdot r_t,$$

where $w_X > 0$ for each $X\in\mathcal{B}$ and $w_t \geq 0$.

  - $(w_X)_{X\in\mathcal{B}}$: weights for the option value of carrying inventory over time for different blood types (long-term forecasts).
  - w_t: social value of transfusing as much blood as possible.

The Blood Allocation Problem

We fix everything except

- the donor profile $D = (D_i)_{i \in I}$ and
- the utility profile $u = (u_i)_{i \in I}$.

A blood allocation problem is represented by a pair (D, u).

Allocations

- Given a problem (D, u), an allocation α is a nonnegative integer vector such that, for every blood type X:

$$\underbrace{\alpha_X(b)}_{\text{leftover}} + \underbrace{\sum_{i:\, X \in \mathcal{C}(\beta_i)} \alpha_X(i)}_{\text{received by patients}} = \underbrace{v_X}_{\text{inv.}} + \underbrace{\sum_{d \in \cup_i D_i:\, \beta_d = X} \alpha(d)}_{\text{supplied by patients}}$$

  and for every patient i, $\sum_{X \in \mathcal{C}(\beta_i)} \alpha_X(i) \leq \overline{n}_i$.

- The schedule of i in α: $\alpha(i) = (\alpha_r(i), \alpha_s(i))$, where $\alpha_r(i) = \sum_{X \in \mathcal{C}(\beta_i)} \alpha_X(i)$, $\alpha_s(i) = \sum_{d \in D_i} \alpha(d)$.
- The set of all allocations for D: $\mathcal{A}(D)$.

Two Policy Levers

1. Feasible Schedule Menus (FSMs)
    - Formalize (potentially endogenous) exchange rates.
    - For each i and reported D_i, the FSM $\mathcal{F}_i$ specifies which schedules (r, s) are admissible.
2. Allocation Mechanisms
    - Pick an allocation given the reports and the FSMs.

Inseparable: FSMs constrain schedules but don't choose; a mechanism cannot ensure proper exchange rates or incentives without operating on a suitable FSM profile.

Feasible Schedule Menus

- $\mathcal{F}_i$: feasible schedule menu (FSM) of patient i, a function that, for each D_i, returns a set $\mathcal{F}_i(D_i) \subseteq \mathcal{W}_i(D_i)$ of feasible schedules, with $\mathcal{W}_i(D_i) = \{0, \ldots, \overline{n}_i\} \times \{0, \ldots, |D_i|\}$.
- For each patient i, a minimum guarantee $\underline{g}_i$ such that for every D_i,

$$\min\{r : (r, s) \in \mathcal{F}_i(D_i)\} = \underline{g}_i \quad \text{or} \quad \mathcal{F}_i(D_i) = \{(0, 0)\}.$$

- Given an FSM profile $\mathcal{F} = (\mathcal{F}_i)_{i \in I}$, restrict attention to

$$\mathcal{A}(\mathcal{F}, D) = \left\{\alpha \in \mathcal{A}(D) : \alpha(i) \in \mathcal{F}_i(D_i),\ \forall i \in I\right\}.$$

Exchange-Rate Policies Around the World

(P1) One-for-one policy (the most common in the world)

$$\mathcal{F}_i(D_i) = \big\{(r, s) \in \mathcal{W}_i(D_i) : s = r\big\}$$

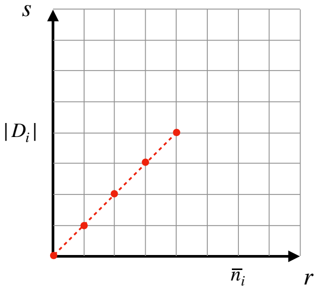

Exchange-Rate Policies Around the World

(P2) Two-for-one policy (Cameroon, Congo, Mexico)

$$\mathcal{F}_i(D_i) = \left\{(r,s) \in \mathcal{W}_i(D_i) : s = 2r\right\}$$

(P3) Fixed donor policy ($x = 1$ in Delhi; $x = 3$ in Turkey)

$$\mathcal{F}_i(D_i) = \begin{cases} \{(0,0)\} & \text{if } |D_i| < x, \\ \left\{(r,s) \in \mathcal{W}_i(D_i) : s = x\right\} & \text{otherwise.} \end{cases}$$

Exchange-Rate Policies Around the World

(P4) Forgiving policy (Tucumán, Argentina): ad hoc, debts may be forgiven

$$\mathcal{F}_i(D_i) = \left\{(r,s) \in \mathcal{W}_i(D_i) : s \leq r \leq |D_i|\right\}$$

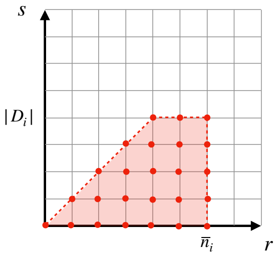

Exchange-Rate Policies Around the World

(P5) Minimum guarantee policy (Xi'an, China): $\underline{g}_i = \min\{3x_i, \overline{n}_i\}$, where x_i is units patient i donated before

$$\mathcal{F}_i(D_i) = \left\{(r,s) \in \mathcal{W}_i(D_i) : s = r - \underline{g}_i\right\}$$

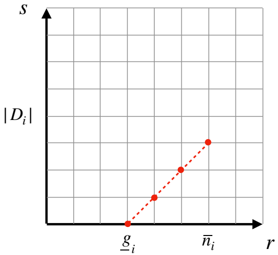

Notice: in P1–P5 (except the ad hoc P4), exchange rates are exogenous.

Mechanisms and Incentives

- Given any FSM profile $\mathcal{F}$, a mechanism f assigns an allocation $f(\mathcal{F}, D, u) \in \mathcal{A}(\mathcal{F}, D)$ to each problem (D, u).
- A mechanism f is incentive compatible under $\mathcal{F}$ if for any (D, u), patient i, $D'_i \subseteq D_i$, and u'_i,

$$u_i\big[\alpha(i)\big] \;\geq\; u_i\big[\alpha'(i)\big],$$

where $\alpha = f\big(\mathcal{F}, (D_i, D_{-i}), (u_i, u_{-i})\big)$ and $\alpha' = f\big(\mathcal{F}, (D'_i, D_{-i}), (u'_i, u_{-i})\big)$.

Mechanisms and Incentives

- A mechanism f is incentive compatible with respect to donors under $\mathcal{F}$ if for any problem (D, u), patient i, and $D'_i \subseteq D_i$,

$$u_i\big[\alpha(i)\big] \;\geq\; u_i\big[\alpha'(i)\big],$$

where $\alpha = f\big(\mathcal{F}, (D_i, D_{-i}), u\big)$ and $\alpha' = f\big(\mathcal{F}, (D'_i, D_{-i}), u\big)$.

Donor revelation is the first-order concern: more reported RDs $\Rightarrow$ more total blood supply $\Rightarrow$ more transfusions.

Mechanisms and Incentives

- A mechanism f is donor monotonic under $\mathcal{F}$ if for any problem (D, u), patient i, and $D'_i \subseteq D_i$,

$$f_r\big(\mathcal{F}, (D_i, D_{-i}), u\big)(i) \;\geq\; f_r\big(\mathcal{F}, (D'_i, D_{-i}), u\big)(i).$$

- A weak and plausible incentive property: blood donation is substantially less invasive than other organ-donation procedures, and blood is replenished quickly.
- Will serve as the foundation for our incentive compatibility results.

Weighted Utilitarian Mechanisms

- A natural class: pick a feasible allocation maximizing the weighted sum of patient utilities and the bank's transfusion/inventory objectives.
- Given positive weights $w = \big((w_i)_{i\in I},\, (w_X)_{X\in\mathcal{B}},\, w_t\big)$, f is a WUM w.r.t. w if

$$f(\mathcal{F}, D, u) \in \arg\max_{\alpha\in\mathcal{A}(\mathcal{F},D)} U(w,u,\alpha), \quad U(w,u,\alpha) = \sum_{i\in I} w_i u_i\big(\alpha(i)\big) + o\big(w, \alpha$$

where $T(\cdot)$ is a small linear one-to-one tie-breaker.

A Special Case of Weighted Utilitarian Mechanisms

- Priority mechanisms: given a priority order over the patients and the blood bank, the priority mechanism lexicographically maximizes their utilities/objectives.
  - A two-way unit-demand priority mechanism for kidney exchange (Roth, Sönmez & Ünver'05, with one-for-one exchange rate).
  - A maximal mechanism with priority tie-breakers for time banks (Andersson, Cseh, Ehlers & Erlanson'21, with one-for-one exchange rate).
  - A priority mechanism for shift exchange (e.g., used in nurse/doctor staffing) (Manjunath & Westkamp'21, with one-for-one exchange rate).

Efficiency

Proposition

Let f be a WUM with respect to w. If the weights $(w_X)_{X \in \mathcal{B}}$ and w_t on the bank's objective are sufficiently small compared to variations in the weighted sum of patient utilities, then f is Pareto efficient.

Intuition: when patients' utilities dominate the objective, the maximizer is on the Pareto frontier of patients.

Fairness

- WUMs can achieve more egalitarian allocations than priority mechanisms by assigning appropriate weights to patients.
- When ρ_i is strictly concave and σ_i is strictly convex, a WUM tends to assign similar schedules to similar patients.

L$^\natural$-convexity

A set $S \subseteq \mathbb{Z}^2_+$ is L$^\natural$-convex (L-natural-convex; Murota'98) if for every $x, y \in S$,

$$\left\lfloor \frac{x+y}{2} \right\rfloor,\ \left\lceil \frac{x+y}{2} \right\rceil \in S.$$

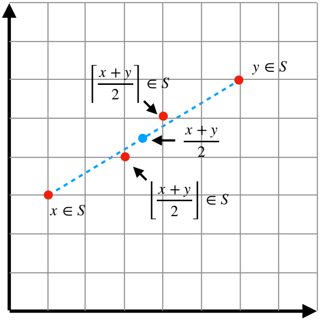

Together with M$^\natural$-convexity, one of the two standard discrete-convexity notions.

$\mathrm{L}^\natural$-convexity

Definition

An FSM profile $\mathcal{F}$ satisfies $\mathrm{L}^\natural$-convexity if $\mathcal{F}_i(D_i)$ is $\mathrm{L}^\natural$-convex for every patient i and her donor set D_i.

L$^\natural$-convexity

Examples:

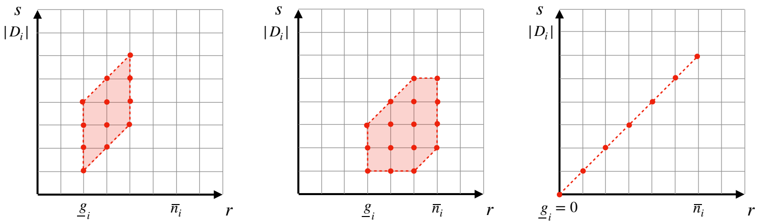

- L$^\natural$-convex feasible schedule set: integral points of an at-most-six-sided convex polygon with side slopes 1, 0, or ∞.

Fairness

Proposition

Consider a WUM f w.r.t. w, an $\mathsf{L}^\natural$*-convex FSM profile $\mathcal{F}$, and a problem (D,u). For any two patients i,j with $\beta_i = \beta_j$, identical compatible-donor type counts, $\rho_i = \rho_j$, $\sigma_i = \sigma_j$, $\mathcal{F}_i(D_i) = \mathcal{F}_j(D_j)$, and $w_i = w_j$:*

- *if ρ_i is strictly concave,* $\left| f_r(\mathcal{F},D,u)(i) - f_r(\mathcal{F},D,u)(j) \right| \leq 1$;
- *if σ_i is strictly convex,* $\left| f_s(\mathcal{F},D,u)(i) - f_s(\mathcal{F},D,u)(j) \right| \leq 1$.

A discrete and approximate "equal treatment of equals."

Incentive Properties

- We start to analyze the incentives under weighted utilitarian mechanisms by considering donor monotonicity.
- This will serve as the foundation for the other incentive compatibility results later.
- For weighted utilitarian mechanisms to be donor monotonic, we need to make 3 regularity assumptions on the feasible schedule menus.

$L^\natural$-convexity

- The first property is $L^\natural$-convexity, introduced earlier for allocative fairness.
- Its key role in establishing donor monotonicity: it prevents hidden complementarities between units of blood received — ruling out holes and rugged edges in feasible schedule sets.
  - Concavity of the blood valuation function ρ_i plays a similar role in incentives as $L^\natural$-convexity of feasible schedule sets.

Feasibility of Positive Price

Definition

A feasible schedule menu profile $\mathcal{F}$ satisfies feasibility of positive price if for every patient i and her donor set D_i, the following holds:

- if $(r,s),(r',s') \in \mathcal{F}_i(D_i)$, $r' > r$ and $s < |D_i|$, then there exists $s'' > s$ such that $(r',s'') \in \mathcal{F}_i(D_i)$; and
- if $(r,s),(r',s') \in \mathcal{F}_i(D_i)$, $r' < r$ and $s > 0$, then there exists $s'' < s$ such that $(r',s'') \in \mathcal{F}_i(D_i)$.

- If the patient can potentially receive a larger (or smaller) amount of blood, then she can potentially receive this amount by supplying more (or less).
- This is the analogue of each unit of blood has a positive “price”.

Feasibility of Positive Price

Examples: both satisfy L$^\natural$-convexity

violates feasibility of positive price

satisfies feasibility of positive price

- A "flat bottom" or a "flat top" is ruled out (whenever possible).
  - They can be flat only at feasibility boundaries, i.e., at $s = 0$ and $s = |D_i|$.

Non-diminishing Favorability in Donors

Definition

A feasible schedule menu profile $\mathcal{F}$ satisfies non-diminishing favorability in donors if for every patient i and pair of her donor sets D_i, D'_i such that $D'_i \subseteq D_i$, we have:

- if $(r, s) \in \mathcal{F}_i(D'_i)$ and $r \geq \underline{g}_i$, then there exists $s' \leq s$ such that $(r, s') \in \mathcal{F}_i(D_i)$; and
- if $(r, s) \in \mathcal{F}_i(D_i)$, $s \leq |D'_i|$ and $(r, s') \in \mathcal{F}_i(D'_i)$, then there exists $s'' \geq s$ such that $(r, s'') \in \mathcal{F}_i(D'_i)$.

- Non-diminishing favorability in donors manifests itself geometrically as $\mathcal{F}_i(D_i)$ being an expansion of $\mathcal{F}_i(D'_i)$ in the direction of receiving more blood or a downward shift of $\mathcal{F}_i(D'_i)$.
- One-for-one exchange satisfies all properties.

An Example Satisfying All Properties

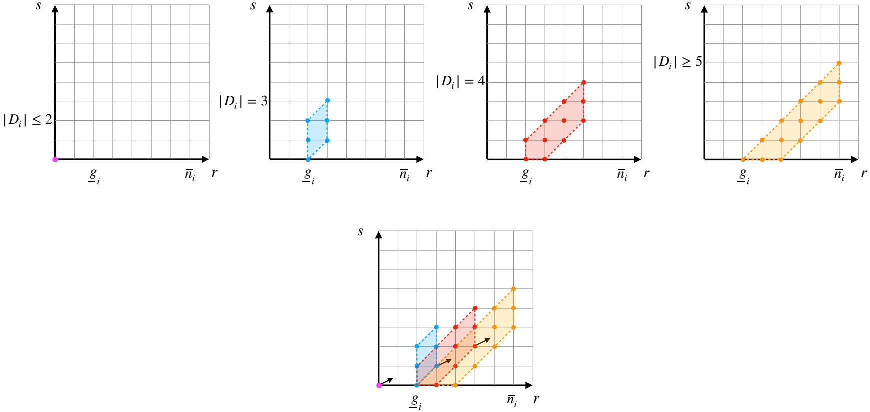

Example: Beyond One-for-one Exchange

- Strict two-for-one exchange (sub-Saharan Africa, Mexico) is not L$^\natural$-convex.
- Instead, use the smallest L$^\natural$-convex envelope satisfying the other properties:

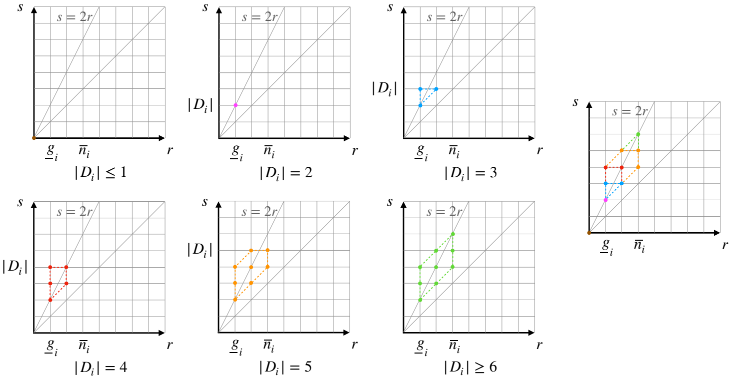

Donor Monotonicity: Main Result

Theorem

Let $\mathcal{F}$ be an FSM profile satisfying $\mathrm{L}^\natural$-convexity, feasibility of positive price, and non-diminishing favorability in donors.

1. *If every patient has a quasi-linear utility ($u_i(r,s) = \rho_i(r) - \theta_i \cdot s$), all WUMs are donor monotonic under $\mathcal{F}$.*
2. *If every patient has a general separable utility, all priority mechanisms (incl. maximal) are donor monotonic under $\mathcal{F}$.*

Two special cases of $\mathcal{F}$ satisfying the three properties: 1. one-for-one exchange 2. whole consumption space

Skip Proof Sketch

Proof Sketch

- Define an auxiliary matching market isomorphic to the original problem.
- A patient is matched with the donors who donate to her and also those of her own donors who do not donate to anyone.
- Consider two extended problems: $D \supset D'$ such that i conceals exactly one of her donors.
- Consider matching M of D and M' of D' s.t. i receives more under M'.

Lemma

There exists a cycle or chain C, from M to M' s.t. each j points to a donor in $M'(j) \setminus M(j)$, each d points to $M^{-1}(d)$ except that in a chain head does not point and tail is not pointed. Moreover, $M' - C$ of D' and $M + C$ of D are well-defined matchings under three conditions of $\mathcal{F}$.

Proof Sketch

- Suppose f is weighted utilitarian with weights w, $f(D) = M$, $f(D') = M'$ and i receives more under M' (recall $D \supset D'$ with only i concealing 1 donor). Let $U(\cdot)$ denote the overall w-weighted sum of utilities.
- Show $M + C$ and M are "welfare" equivalent (for quasi-linearity):
  1. Suppose not: then $U(M) > U(M + C)$.
  2. A patient receives or supplies one more (less) unit under $M + C$ wrt $M \iff$ the same happens under M' wrt $M' - C$ (Lemma). By concavity of value functions and overall quasi-linearity of utility functions, $U(M) - U(M + C) \leq U(M' - C) - U(M')$.
  3. Therefore, $U(M' - C) > U(M')$, where $M' = f(D')$ and $M' - C$ is a feasible matching for D', a contradiction.
- Repeat cycle/chain addition to $M + C$, welfare equivalence continues. Every repetition ensures we get closer to M'.
- Eventually, we reach M' and then M and M' should be welfare equivalent, a contradiction $\Rightarrow f$ is donor monotonic. QED

Individual Rationality

Definition

A feasible schedule menu profile $\mathcal{F}$ satisfies individual rationality if for every patient i and her donor set D_i, we have $r \geq s$ for any $(r, s) \in \mathcal{F}_i(D_i)$.

Strong Non-diminishing Favorability in Donors

Definition

A feasible schedule menu profile $\mathcal{F}$ satisfies strong non-diminishing favorability in donors if for every patient i and pair of donor sets D_i, D_i' such that $D_i' \subseteq D_i$, we have:

- if $(r, s') \in \mathcal{F}_i(D_i')$ and $r \geq \underline{g}_i$, then there exists s such that $(r, s) \in \mathcal{F}_i(D_i)$; and
- if $(r, s) \in \mathcal{F}_i(D_i)$ and $(r, s') \in \mathcal{F}_i(D_i')$, then $s \leq s'$.

Incentive Compatibility w.r.t. Donors: Main Result

Theorem

Suppose every patient has a general separable utility function. If a feasible schedule menu profile $\mathcal{F}$ satisfies

- $\mathrm{L}^\natural$*-convexity,*
- *feasibility of positive price,*
- *strong non-diminishing favorability in donors,*
- *individual rationality*

then every WUM is incentive compatible with respect to donors under $\mathcal{F}$.

Incentive Compatibility: Main Result

Theorem

Suppose every patient has a general separable utility function. If an FSM profile $\mathcal{F}$ satisfies

- $\mathrm{L}^{\natural}$*-convexity,*
- *feasibility of positive price,*
- *strong non-diminishing favorability in donors,*
- *individual rationality,*

then all priority mechanisms (including maximal mechanisms) are (fully) incentive compatible under $\mathcal{F}$.

Skip Policy

Simple Menu Design

- Under L$^\natural$-convexity and feasibility of positive price, every menu can be written as

$$\mathcal{F}_i(D_i) = \Big\{(r,s) : \underline{r}_i \le r \le \overline{r}_i,\ \max\{r-\ell_i,0\} \le s \le \min\{r+k_i,|D_i|\}\Big\}.$$

- A simple design: fix $\underline{g}_i, \overline{n}_i, k_i, \ell_i$; set $\underline{r}_i = \underline{g}_i$, $\overline{r}_i = \overline{n}_i$ for all D_i. Satisfies all three properties.

- Easy-to-explain rule:

  "You receive at least $\underline{g}_i$ and at most $\overline{n}_i$ units; you supply at most k_i more, and at least ℓ_i less, than what you receive."

Choosing the Mechanism

- Priority mechanisms (incl. maximal) are natural: incentive robustness, easy to explain, arrival-time priority minimally departs from FCFS.
- When preferences are nearly degenerate: utility design can be used as an alternative where designer picks ρ_i, σ_i instead of eliciting them.
- When utilities are hard to elicit: the maximal mechanism maximizes total transfusions.

Fairness via Endogenous Exchange Rates

- Equitable FSM: a patient always receives at least $\underline{g}_i$, even with no donors; she reaches $\overline{n}_i$ as donors expand.
- Compatible with practice (e.g., Cambodia waives RDs for patients with no next-of-kin).

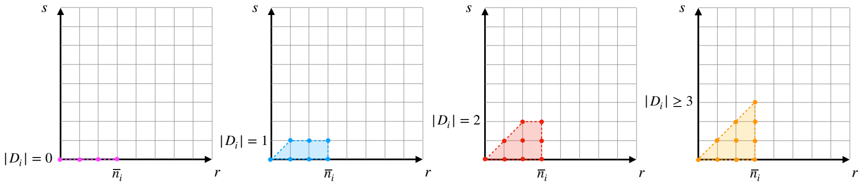

Blood Quality

- FSMs can depend on observable donor characteristics — repeat donors carry the least infection risk.
- Menus rewarding repeat-donor revelation $\Rightarrow$ concealing one strictly lowers utility.

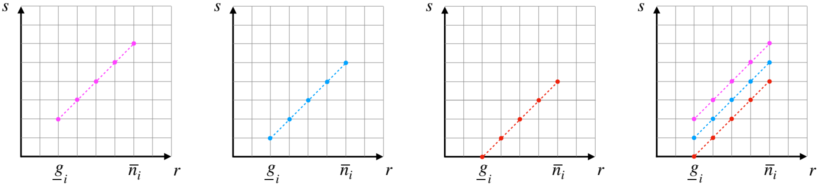

# of repeat donors = 0   # of repeat donors = 1   # of repeat donors $\geq$ 2

Hybrid Implementation with VNRDs

- Two policy variables:
  - Non-emergency patients: simple four-parameter FSM (the typical RD case).
  - Emergency patients: forgiving menus + immediate service from inventory.
- Mechanism: a priority mechanism — emergencies prioritized over non-emergencies; blood bank last.
- Timeline:
  1. Serve emergencies dynamically from inventory; if not enough, bundle with non-emergency batch giving emergencies top priority.
  2. Non-emergencies submit requests + RDs; bank runs the mechanism daily; donors donate; blood is distributed.

Simulations: Setup

- Patient population $|I| \in \{25, 100\}$; 1,000 random markets per scenario.
- Donor number $|D_i| \in \{0, 1, \ldots, 6\}$ uniform iid.
- Maximum need $\overline{n}_i \in \{1, \ldots, 6\}$ (± 1) uniform iid.
- Inventory ratio $\iota \in \{0,\ \frac{1}{50},\ \frac{1}{25},\ \frac{1}{10},\ \frac{1}{5},\ \frac{1}{2},\ 1\}$ of total max need.
- Blood-type distribution: India.
- Three mechanisms under one-for-one exchange: FCFS; priority (same order as FCFS, bank last); the FCFS-dominating IC mechanism.

Simulations: Total Transfusions

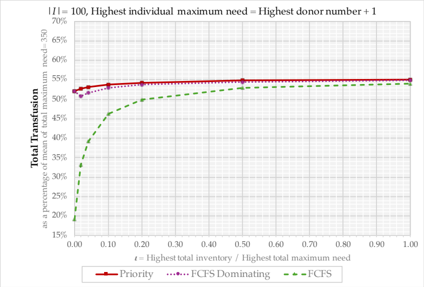

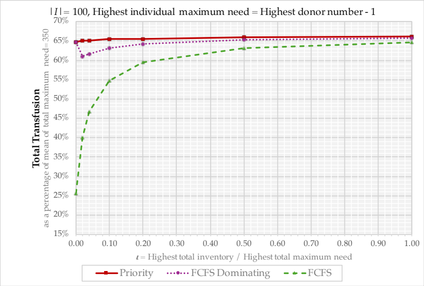

Simulations: Main Takeaways

1. Low inventory requires exchange among patients. At $\iota = 0$, FCFS attains only $37 - 44\%$ of priority's transfusions.
2. FCFS approximates efficiency at moderate inventory. At $\iota = 0.5$, FCFS reaches $\sim 96\%$ of priority.
3. Priority is least sensitive to ι; FCFS is most sensitive.
4. Most patients prefer priority to FCFS. At $\iota = 0$: $54 - 65\%$ strictly prefer priority; share preferring FCFS never exceeds 1.9%.

Extension: General Multi-Unit Exchange

- Suppose compatibility relations $\mathcal{C}_i$ are also private information.
- Full incentive compatibility $=$ truthful revelation of $\mathcal{C}_i$, u_i, and D_i.

Proposition

Under one-for-one exchange, all priority mechanisms (incl. maximal) are fully incentive compatible.

- Applications: time banks (Andersson et al.'21); shift exchange (Manjunath & Westkamp'21).

Conclusions

- A novel framework for multi-unit exchange with flexible exchange rates; leading application is blood allocation.
- FSMs + WUMs achieve efficiency, approximate equal treatment of equals, and incentive compatibility under $L^\natural$-convex menus plus two natural conditions.
- Practical design: four-parameter menus, hybrid implementation with VNRDs, fairness and blood-quality designs.
- In the spirit of minimalist market design (Sönmez'24): improving the RD system can address chronic shortages where VNRD systems alone have not.
- Beyond blood: time banks, shift exchange, and other multi-unit exchanges.