

A model for n -trading in ongoing markets in heterogeneous indivisible goods

Bruce Duggan

Abstract

A variant of the assignment problem is used to model ongoing markets designed to facilitate trades of heterogeneous indivisible goods which potentially involve more than two participants in each trade (n -trading). A process for making comparisons between modelled markets and their rules is outlined. The results of a computer-simulated comparison between traditional two-sided trading, n -trading, and the primary alternative to two-sided trading in the game-theoretic literature (Gale's top-trading cycle procedure developed by Shapley and Scarf; 1974) are reported. They indicate that agents in the modelled ongoing market have a preference for both n -trading and Gale's procedure over traditional two-sided markets, and patient agents have a preference for n -trading over the Gale procedure. A second computer simulation comparing n -trading with two-sided trading in ongoing markets ranging in size from 2 to 35 agents indicate that n -trading results in greater benefits for agents than two-sided trading. This difference holds over all market sizes. As well, the results of a prototype test suggest that real-world n -trading is practical. Applications may be quite broad, including trading in fisheries and resource permits under ocean zoning regimes.

Keywords

n -trading; heterogeneous indivisible goods; market preference; game preference; assignment problem; assignment games; cooperative game theory; market design; mechanism design.

1 Introduction

This paper develops an optimization mechanism (n -trading) for *trading*¹ in heterogeneous indivisible *goods*² in ongoing *markets*³.

To place n -trading in context, the paper begins by developing a method for making comparisons between models of markets that draws on market design approaches, and reviews existing models for optimizing trade in heterogeneous indivisible goods. The general optimization formulation for n -trading is given, and then compared to the existing models from operations research and from cooperative game theory. This is followed by a report on the results of a computer simulation comparing n -trading to traditional two-sided trading, as well as the results of a test of the practicality of the n -trading approach in a prototype market. Finally, potential applications are outlined.

¹ Trading is defined in section 2.5. Defined terms are written in italics the first time they are used.

² A good is anything to which at least one agent is able to assign a utility function. A good can be individual or bundled, divisible or not. This definition includes goods, services (including work), and money. It also includes other agents (in, for example, a marriage market). This paper focuses only on the subset of indivisible goods. Further qualities of these goods are given in the Appendix.

³ Markets are defined in section 2.5. Further qualities of modelled markets are given in the Appendix.

2 Making comparisons

2.1 Market preferences

We routinely compare markets as we make choices about which markets to participate in. A public corporation must consider which stock markets should list their stocks; a person in the market for a car weighs the merits of buying online vs. going to a dealership. Understanding why and how potential participants choose one market over others is not a simple matter. One promising approach is the emerging field of market design. Market design is both a very old field and a very new one. Plato (*ca.* 360 BCE) devotes part of Book II of his *Republic* to an exercise very close to market design, positing the ideal market for the ideal state. Yet, today, it is still possible for one of its leading practitioners to consider it an emerging discipline (Roth, 2002). (See also Niederle *et al.*, 2007, for an overview.) To date, market design has not focused much on market-choice behaviour. However, the methodologies being developed by it can be usefully applied to model market choice.

2.2 Game preferences

In making them “rational”, von Neumann and Morgenstern created game players compelled to always choose utility maximization, treating game rules as “absolute commands” (Von Neumann and Morgenstern, 1944: 49). Players do not truly exercise choice; to say a particular player has a “preference” is simply a shorthand way of saying that one (pure or mixed) strategy will (certainly or probably) provide a higher payoff than others for that player. This approach was extremely useful in constructing precise models. However, it may have hampered the modelling of markets in which not all rules are always followed (which is undoubtedly the majority of actual markets). It also meant that issues of game comparison, preference, and modification received little attention over the initial years of game theory’s development.

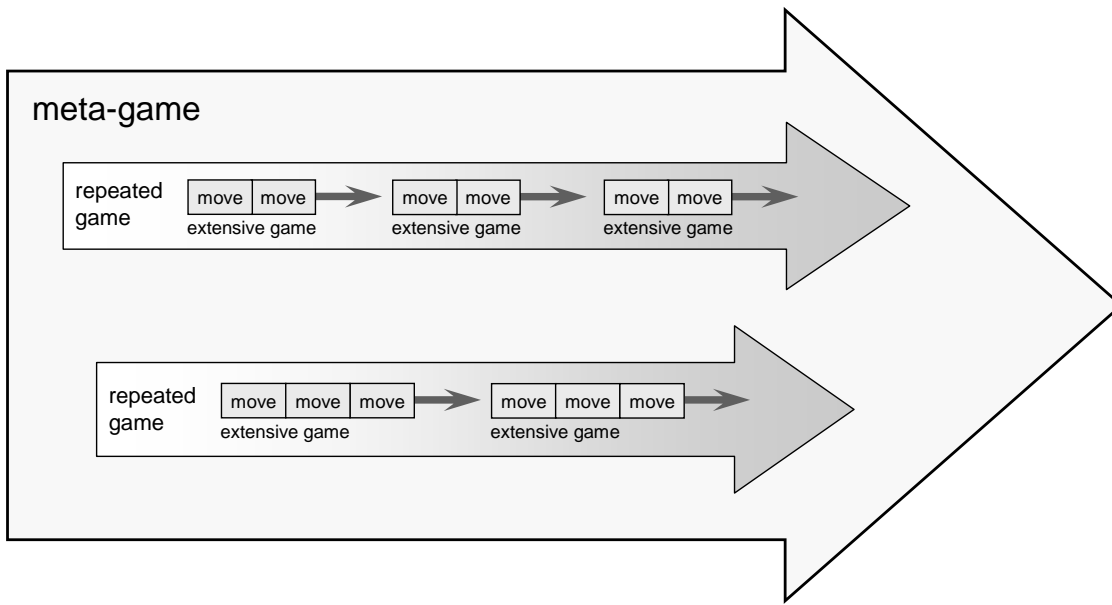
Shapley (1953) may have been the first to suggest that, in addition to having preferences within games, players also have preferences over games.⁴

If given a choice between games, players will prefer the game that (certainly or probably) yields a higher payoff for them. This situation could be modelled as a single extensive game, making the choice of one sub-game over available alternative sub-games part of the sequence of actions. However because, in actual markets, market-choice behavior occurs prior to, simultaneously with, and after in-market behaviours, it is more useful to model choosing between markets as a meta-game.⁵ A meta-game model of ongoing markets will have four levels to consider—strategic, extensive, repeated, and meta:

Figure 1. Framework for modelling market-choice behaviour between two ongoing markets.

⁴ The idea was implied, rather than fully developed; Shapley pointed out that players would have preferences over whether or not to play a game.

⁵ The meaning here is similar to, but not identical with, Howard’s metagame (1971).



Game preferences arise in the meta-game. A player in a meta-game can be expected to compare vNM utilities between games prior to participating in them. As they participate, they will compare game experiences (including actions of other players, completeness of information, available payoffs, and options for withdrawal). After participating, they will evaluate the differences in payoffs between games, as well as the distribution of those payoffs between players. They may evaluate the game in isolation (a one-person meta-game) or while considering the evaluations and actions of other players within the meta-game. Just as a player may have a pure or a mixed strategy within a game, they may have a pure or mixed game-preferences in the meta-game.

In using games to model markets, we follow Shapley's (1955) lead and model them as cooperative games.⁶ A cooperative game is unstable or stable, depending on whether or not players have an incentive to defect from the coalition. Similarly, in a meta-game, a repeated cooperative game can be *unsustainable*⁷ or *sustainable*, depending on whether players have an incentive to defect from the game.

However, there is a crucial difference between stability and sustainability: Stability can be determined by considering only the game itself. Not so with sustainability; a game is

⁶ Expanded elsewhere, including in Shapley and Shubik (1969).

⁷ A repeated game is unsustainable if the payoffs for players are less than the payoffs to an available alternative game. An ongoing market becomes unsustainable if agents stop participating in it, even through the rules do not require that the market terminate at that point. (The minimum number of agents required for a sustainable market is 2.) In some literature, "unstable" or "subject to unravelling" is used instead of "unsustainable" (e.g.: Roth, 2000). Because "unstable" has a particular meaning within cooperative games, that term is not used here. "Unravelling" is not used here because it has no modern antonym.

sustainable only if available alternatives games have lower payoffs.⁸ This view of sustainability has a number of implications:

- A repeated game, even if it has stable outcomes at each stage, is not necessarily sustainable.
- The corollary is also true: A repeated game with unstable stage outcomes may nonetheless be sustainable.
- A sustainable game can become unsustainable if a new one becomes available, even through the rules of the first game have not changed. If players can receive higher payoffs in the new game, they will desert the old one.
- A player may remain in a game, even if another has higher payoffs, depending on the cost of switching. Game-switching costs serve as a counterbalance to game-preferences.
- A player who discounts⁹ ($\delta > 0$) will also consider the timing of costs compared to benefits. Players will weigh the discounted cost of switching against the relative discounted *benefit*¹⁰ of the alternative game(s). A number of factors will weigh in this decision: the player's discount rate, the timing of switching costs and benefits, and the probabilities of those costs and benefits.

These implications suggest that a meta-game modelling of markets may be useful, because many of the phenomena in the meta-game (the disruptive effects of new games/markets, the role of switching costs in sustaining less beneficial games/markets, the role that timing plays in assessing both the benefits and costs of the games/markets being compared, the rationality of participating in more than one game/market at a time) may plausibly have counterparts in actual markets. The meta-game model also enables us to model market comparison, preference, and modification without giving up the clarity of Von Neumann and Morgenstern's inexorably utility-maximizing player.

2.3 *Rules preferences*

What are the rules which men naturally observe in exchanging them [goods] either for money or for one another, I shall now proceed to examine. These

⁸ Depending on the games, the higher payoff can be either certain or probable. A game with a higher payoff is weakly sustainable if the alternatives available have payoffs equal to it; it is strongly sustainable if all available alternatives have lower payoffs.

⁹ An impatient agent has a high discount rate—strongly preferring benefits now to benefits later. A patient agent has a low discount rate—valuing benefits later almost as highly as benefits now. An infinitely patient agent ($\delta = 0$) values present and future benefits equally. An infinitely impatient agent ($\delta = 1$) has the same incentives as an agent who will participate in only one market round.

¹⁰ Benefit is defined in section 2.5.

rules determine what may be called the relative or exchangeable value of goods.¹¹

Adam Smith (1776)

No matter how free, no ongoing market is chaotic.¹² Markets have a nearly infinite variety of mechanisms governing them. Some are unwritten (customs, norms, rules) and may not even be consciously understood by all market participants.¹³ Some (regulations, laws) are explicit, can be extraordinarily complex, and often include sanctions for those who violate them. We will call the class of all mechanisms governing all markets *Rules* (\mathcal{R}).¹⁴ Because they are present in one form or another in all ongoing markets, we can classify those markets by the rules that govern their functioning.

We can expect market participants to prefer some rules over others. We might expect a corporation to prefer to list their shares on a stock market with simple prospectus requirements over one with complex requirements. We might expect someone buying a car to prefer a rule requiring sellers to disclose all known defects over one that allowed sellers to withhold information. Understanding why participants prefer one set of rules over another is every bit as complex as understanding how and why they choose one market over another. Fortunately, one of the fundamental principles of game theory enables us to model rule preference.

2.4 The relationship between market preferences, game preferences, rule preferences, and market design

As von Neumann and Morgenstern defined games, a “game is simply the totality of the rules which describe it” (*op. cit.*). As a result, a preference between games is identical with a preference between rules.

¹¹ *The Wealth of Nations* 1 (4) para. 12.

¹² Hayek (1948) went so far as to argue that market rules arise “spontaneously” out of human interaction. Holding the opposite view, Keynes and the architects of trade regimes flowing from The Bretton Woods Agreements (1944) believed markets require extensive, explicit rules.

¹³ The past 50 years have seen the beginnings of an understanding how rules and norms are constructed, and how they change. Gardner and Ostrom (1991) and Ostrom *et al.* (1994) are notable, particularly because they build on Von Wright’s rigorous analysis of norms (1963), beginning an integration of his work with game theory. They consider the relationship between games, rules, and norms both theoretically and in applied contexts. Other work exploring norms from a game theoretic perspective include Kandori (1992) and the work he cites, as well as Hasker (2007). Considerable work is also being done for a sociological perspective on the role of norms in markets (Biggart and Beamish, 2003, provide a summary.) As well, the centrality of trust in social institutions (including markets) has received increasing attention in recent years. Misztal (1996) provides a comprehensive sociological frame; and Castaldo (2008) integrates research on the role of trust in market functioning. To say a participant has “trust” in a market encounter is to say they have two beliefs: that trading in keeping with that market’s rules will probably benefit them, and that others involved in that encounter will follow the rules often enough for them to realize that benefit.

¹⁴ The mechanisms governing a particular class of markets (for example, all markets facilitating two-sided trades) will also be called *rules* (R). The mechanism(s) governing a single market will be its *ruleset* (r).

If game X produces higher payoffs for a player than game Y , they will prefer X -rules over Y -rules. This preference can exist even if game X does not (yet):

...when the outcomes are unstable, agents have incentives to *change* the rules of the game, as when they decide to introduce a centralized matching procedure, or to defect from one (Roth and Sotomayor, 1990: 243).

Perhaps market design can be best characterized as the comparative study of market rules. As a theoretical field, it explores what rules *agents*¹⁵ prefer in modelled markets. As an applied field, it seeks to create rules for actual markets which participants should prefer (if they behave, at least some of the time, like utility-maximizing agents).

2.5 Defining markets

So far, we have used the term “market” without a definition. We need at least an informal one.

A market as an institution in which the owners of goods trade, with the aim of benefiting themselves.¹⁶ This definition may be useful in characterizing actual markets. However, to be able to compare modelled markets (and their rules) more precision is needed, particularly regarding what a benefit is, what a trade is, and what process agents go through when they trade.

2.5.1 Benefits

Let the utility function of good h_j to agent a_i be v_{ij} . If a_i enters a market round¹⁷ with h_i and leaves with good h_j (after trading h_i for h_j), then a_i 's *benefit* (b_i) is:

$$b_i = v_{ij} - v_{ii}.$$

The total benefit of a market round (B) is:

$$B = \sum_{i=1}^n b_i \text{ where } n = |pA|, \text{ the number of potential traders in that market.}$$

The total benefit (\mathcal{B}) of a market is:

$$\mathcal{B} = \sum_{i=1}^p B_i \text{ where } p = |D|, \text{ the number of rounds in that market.}$$

¹⁵ An agent is a participant in a modelled market. An agent is the equivalent of a player in a game, and a trader in an actual market. Further qualities of modelled markets are given in the Appendix.

¹⁶ Of course, many other definitions of a market are possible.

¹⁷ A market round will be defined in the next section.

The *benefit ratio* or *benefit percentage* (β) is:

$$\beta_i = \frac{U_{ij} - U_{ii}}{U_{ii}}$$

2.5.2 Trades

In the markets modelled here, “*to trade*” is to agree that the ownership rights over one good have been surrendered in exchange for the ownership rights of another good. In these markets, ownership rights include the right to use, consume and trade that good—rights enabling an agent to realize the utility of the good.¹⁸

In the simplest version of this model, a market consists only of a single trade between only two agents, each owning only one indivisible good.¹⁹ In this simplest version, trades are costless. After they examine each others’ wares, each *judges*²⁰ their own good to be preferable to, equal to, or less preferable than the other’s good. So, even in this simplest of markets, nine preference profiles are possible:

Table 1. Preference profile: Relative utility.

		Trader 2		
		prefers own good to other’s good	indifferent between goods	prefers other’s good to own
Trader 1	prefers own good to other’s good	$v_2 < v_1, v_1 < v_2$	$v_2 < v_1, v_1 = v_2$	$v_2 < v_1, v_1 > v_2$
	indifferent between goods	$v_2 = v_1, v_1 < v_2$	$v_2 = v_1, v_1 = v_2$	$v_2 = v_1, v_1 > v_2$
	prefers other’s good to own	$v_2 > v_1, v_1 < v_2$	$v_2 > v_1, v_1 = v_2$	$v_2 > v_1, v_1 > v_2$

Only one preference profile—the bottom right cell—has the potential for a trade in which both participants benefit. In that preference profile, we commonly say both are “motivated to trade” and we expect they will “agree to trade”. If one of the two goods traded is money, we say we have a “willing buyer” and a “willing seller” and the two should be able to “agree on a price”. These surface agreements are built on a crucial

¹⁸ Trading, then, is a social agreement. To *own* something is to have a cluster of rights over it. To characterize trading as an exchange of rights rather than an exchange of goods may appear overly precise. However, it makes the inclusion of services within the definition of goods easier. Most definitions of trading make the inclusion of marriage markets in trading models difficult, and this definition is no exception. (The difficulty arises because the agent and the good are identical.) However, because marriage markets have been central to the understanding of assignment games (discussed below) they need to be included. The difficulty is more apparent than real: In an unmarried state I have exclusive rights over my self and, when I marry, I surrender at least some portion of those rights, and gain a portion of rights over the other person in the marriage. (Of course, one of the rights I do not gain is the right to trade.)


¹⁹ In this simplest version, no other agents or influences on these two agents exist. Questions of who makes an offer first, how that offer is conveyed, and how it is responded to, how agreement is reached, and how the exchange of goods is actually made, while relevant in actual markets, are not dealt with here.

²⁰ To *judge* is to assign a utility function which is, at minimum, ordinal, to the cluster of rights over a good.

underlying disagreement—a disagreement on the relative utility of goods. This disagreement makes trading within that preference profile a positive-sum encounter. It is this disagreement, replicated over an incalculable number of encounters, that creates and sustains all markets, from the smallest informal swap to the largest international system.


Let us make only one change to our simplest of markets: We add a third trader, with an indivisible good of their own. In virtually all actual markets, if Trader 3 wants to trade with Trader 1, he has only three options: pre-empt 2’s offer, convince 1 to reject 2’s offer, or convince 2 to withdraw. This limitation to two-sided (bilateral) trading holds in almost all existing markets.

Figure 2. A bilateral trade.



However, there is no insurmountable reason why a trade must be bilateral, either in actual or modelled markets. The simplest trade involving more than two participants would be:

Figure 3. A trilateral trade.



Actual three-way trades occur occasionally (*e.g.*: Fine, 2008), but they are *ad hoc* and rare. The practical limitations that have constrained almost all trades to two participants include complexity and trust constraints, research costs in finding willing participants to a trade, and transaction costs in completing trades involved more than two participants. Perhaps only one model for more-than-bilateral trading exists in the literature—the Top Trading Cycle²¹—invented by David Gale and reported in Shapley and Scarf (1974).

²¹ “Cycle” was their term for the members of a set of traders who exchange goods between themselves in a particular pattern. A cycle could also be called a trading loop, circle, or ring. Because of the importance of the Shapley-Scarf paper, “cycle” will be used here. A 2-cycle is the set of two traders in a familiar bilateral trade. A 3-cycle is a set of three traders who each “give” their good to the next member in the cycle. (Note that a 4-cycle is a set of four traders who each hand their good to the next member in the cycle: $i \rightarrow j \rightarrow k \rightarrow l \rightarrow i$. If i and j exchanged goods, and k and l exchanged goods, this would not be a 4-cycle, even if the exchange happened only because of an agreement between all four traders; it would be a pair of 2-cycles.)

Outlining *what* a trade is may be simple. We must also outline *how* a trade is done, which is less simple.

2.5.3 Process

All market models compared in this paper will have the same process: a series of steps taken by agents. These steps are sequential in time, and can be imagined as occurring in a physical space:

Step 1: A set of *agents* $A = \{a_1, \dots, a_n\}$ each bring one corresponding heterogeneous indivisible good $H = \{h_1, \dots, h_n\}$ to a trading venue.

An agent taking this step is a *potential* trader (${}_p a$).

Market size = $|{}_p A|$. We call a market with four potential traders is a “4-agent market”.

Step 2: All potential traders assign a utility function to the goods of all potential traders.

If a potential trader finds one or more goods on offer which they judge to be *significantly better*²² than their endowed (original) good, they are *motivated* to trade; otherwise they are *unmotivated*. Based on these utility functions, every element of the set ${}_p A$ is either a *motivated* trader (${}_m a$) or an *unmotivated* trader (${}_{-m} a$).

Iff for any ${}_p a_i$, any $v_{i-i} - t_i - v_{ii} > 0$, then ${}_p a_i \rightarrow {}_m a_i$, otherwise ${}_p a_i \rightarrow {}_{-m} a_i$, where v_{i-i} is the utility to agent i of any good h not brought to the market by agent i .

States: Two states are possible after this step: Either (ε) less than two traders are motivated, or (ε') at least two of the traders are motivated.

Step 3a: If less than two traders are motivated, no trades are *feasible*,²³ the market round ends, and everyone leaves with their original good.

If ε , where $|{}_m a| < 2$, then each element of the set ${}_{-m} A$ forms a 1-cycle.

The market round is *unviable*.

A 1-cycle is a set of one potential trader (a singleton) who comes to a market, considers trading, decides not to, and then leaves with the good they came with.

²² Good h_j is significantly better than good h_i if $v_{ij} - t_i - v_{ii} \geq 0$. The benefit of trading must match or exceed the cost of trading.

²³ In the literature of assignment games and top trading cycles, a “feasible trade” is any trade which conforms to the rules governing those markets. Because of the rules of those markets, all feasible trades are also beneficial to the agents involved. The same meaning is used here.

➤ Ending 1: The step series ends.

Step 3b: If at least two of the traders are motivated, goods are allotted according to the rules of that market.

If ε' , where $|{}_m a| \geq 2$, the market is *viable*.

Apply the *ruleset* (r) to the market.

States: Two states are possible after this step: Either (κ) no traders can find a feasible trade, or (κ'), at least two traders can find such a trade.

Step 4a: Even if they are motivated, unless at least two traders can find a feasible trade, no trades are possible, the market round ends, and everyone leaves with their original good.

If κ , where no n -cycles exist in which $n \geq 2$ (no trades are feasible). Each element of the set ${}_m A$ forms a 1-cycle. All agents are *unsuccessful* traders ($_{-s} a$).

The market round is *unsuccessful*.

➤ Ending 2: The step series ends.

Step 4b: Those who can, trade.

If κ' , each agent i able to join an n -cycle where $n \geq 2$ surrenders their good h_i , incurs a *trading cost*²⁴ (t_i), receives good h_{-i} , and is a *successful* trader ($_{s} a$).

Those who can't trade, don't.

Each element of the set ${}_m A$ not an element of the set ${}_s A$ forms a 1-cycle.

All traders who have not been able make a feasible trade leave with their original good. All traders who have made a trade leave with their new good.

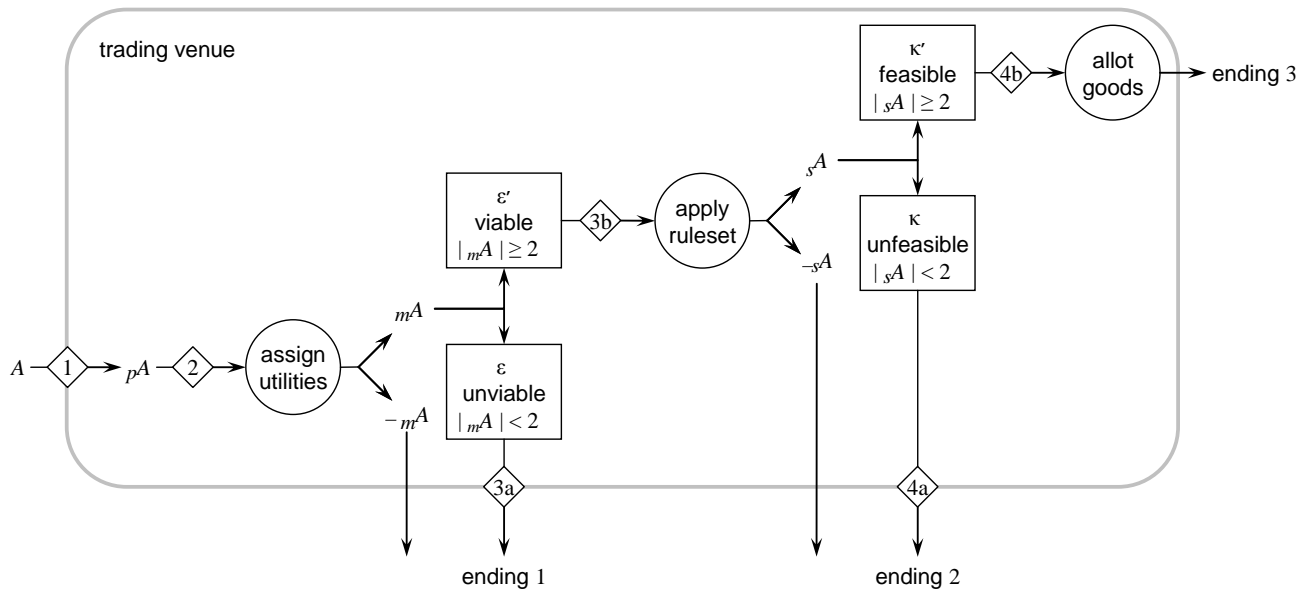
The market round is *successful*.

²⁴ In many market models, the act of trading is costless (the act of trading has a utility = 0). Although this simplifies calculations, it abstracts the model away from what is being modelled, making it less useful. In virtually all real-world markets, actions by participants have a non-zero utility. In most cases, the utility is negative, because actions require effort. Unless stated otherwise, in all markets modelled in this paper, the act of trading has a cost (utility < 0). Agent i incurs a trading cost t_i when they make a trade; it is not incurred if the agent leaves with the same good they arrived with. In other words: Trading takes effort, but it doesn't cost anything to look. The model could be easily adapted to include costs commonly incurred in actual markets (costs for looking, making an offer, accepting an offer, *etc.*).

➤ Ending 3: The step series ends.

This series of steps, completed in order, is a *modelled market round* or *round* (d), analogous to a stage game. A round can occur whether or not trades occur, and even if no agents are motivated to trade. All that is necessary is for the step series outlined above to reach one of its three possible ends.

Figure 4. Trading process.



One or more market rounds, sequential in time, is a *modelled market* (m). The number of rounds can be finite (≥ 1), indeterminate, or infinite.

There are a number of implications of characterizing a market this way:

- If the rules change, the market ends and a new one begins, whereas game theory may well treat that situation as a dynamic game.
- A market can continue even if the agents are not the same in each round, provided the rules do not change and the number of potential traders (p^a) is the same in each round.

What is called a model market here is not exactly what game theory calls a repeated extensive game, although it is very similar.

If markets could be designed that facilitated trading cycles of any size, and if the benefits to the participants in these markets exceeded the benefits of traditional bilateral markets, both market participants and market facilitators should prefer those markets over traditional ones.

3 Existing models for markets in indivisible goods

Three main models for trading in heterogeneous indivisible goods have been developed in the literature—the assignment problem (from operations research), the assignment game, and the top trading cycle (the last two both from cooperative game theory).

3.1 Optimization: Assignment tools

Linear (and, later, quadratic and non-linear) optimization has generally treated the allocation²⁵ of indivisible goods under the rubric of the “assignment problem”. The allocation of goods is, of course, not the same as the trading of goods. However, the techniques developed in operations research have proved adaptable to trading applications.

In the prototypical assignment problem, a decision-maker is given the task of matching a number of workers to an equal number of machines. The workers vary in their efficiency on the various machines, and so the decision-makers’ goal is to optimally assign workers to machines to maximize profit (or minimize costs). The optimization tools used for the assignment problem were initially developed to address what were called shipment problems.²⁶ In the prototypical shipment problem, a decision-maker’s task is to ship goods from supply depots to demand depots, and their goal is to use the minimum transportation resources (or costs) possible. Optimization tools for these shipment problems were primarily developed by Hitchcock (1941) and Koopmans (1947). These approaches were some of the first linear programming mechanisms to be used in industry (Smith, Jr., 1955). Votaw and Orden (1952) recognized the assignment problem as a particular case of the shipment problem and Kuhn (1955) invented an algorithm for solving assignment problems. (He called it the “Hungarian Method”, because he based it on the earlier work of Hungarian mathematicians König and Egerváry).

Koopmans and Beckmann (1957) highlighted the value of these tools in ensuring the optimal allocation of indivisible resources. They outlined their immediate practical use to industry, using the optimal assignment of industrial plants to locations as an example. Perhaps more intriguingly, they suggested that assignment tools might be relevant to

²⁵ Koopmans and Beckmann (1957) used the term “allocation” in connection with markets. Some later literature (e.g.: Abdulkadiroğlu and Sönmez, 1998; Abraham *et al.*, 2004) have reserved “allocation” to refer only to assignment mechanisms where goods are not privately owned. Others (e.g.: Miyagawa, 2002; Niederle *et al.*, 2007; Quint, 1997; Pápai, 2003; Svensson, 1999) do not make this distinction. In this paper, when goods are not privately owned, *initial arrangement* will be used to describe the situation at the start of an assignment problem, while *allocation* will be used to describe the situation after the assignment tool is applied. *Endowment* will be used to describe the ownership of goods prior to a market encounter, while *allotment* will be used to describe the ownership of goods after. (In a 1-cycle, the allotment is the same as the endowment; in all other cycles, the endowment and allotment will be different.) *Start* will be used to include both initial arrangements and endowments. *Assignment* will be used to include both allocations and allotments.

²⁶ Particularly in the early literature, “shipment” problems are sometimes called “transshipment” or “transportation” problems (e.g.: Orden, 1956). The three terms were usually synonymous, and “shipment” is used here to cover all three terms.

larger economic questions, such as the allotment of indivisible goods in competitive markets.²⁷

In the years that followed these initial explorations, a growing array of assignment tools entered the toolkit of operations research, and assignment tools were integrated into commercial optimization software packages. Assignment tools were used to tackle problems as diverse as mobilizing officers in wartime (Bausch, *et al.*, 1991), and reducing power consumption in mobile computers (Yu and Prasanna, 2003). Öncan (2007) provides an overview of various assignment tools, and some of the problems to which they've been applied.

In almost all cases, underneath this diversity remains the fundamental approach developed by Koopmans and others: the use of optimization mathematics to reveal the optimal matching of elements, from two different sets, into pairs.

3.2 Game theory: Assignment games

Game theory has also considered the allotment of indivisible goods, developing ideas parallel to, and often intertwined with, the development linear programming. Von Neumann (1953) may have been the first to publish an equivalence between assignment tools and a proposed game—demonstrating that the specific assignment tool he used and the game he proposed produced identical outcomes.

Other game theory writers continued the idea that assignment tools and games were equivalent, but replaced Von Neumann's game with others. The seminal work was Gale and Shapley's (1962) proposals for producing stable marriages (and other stable pairings). The light-hearted tone disguises a serious purpose: to develop game-theoretic approaches to what have come to be called two-sided matching problems. In the prototypical example, men and women are matched into pairs, creating stable (heterosexual) marriages. Shapley and Shubik (1972) extended this into what they called "the assignment game" to model two-sided markets (the two sides in their game being buyers and sellers), trading indivisible goods for money. They demonstrated that their assignment game was equivalent to a particular assignment problem.

Variants of the assignment game have been used to model a wide variety of situations of interest in economics. Roth and Sotomayor (1990) provide an overview of assignment game variants, and some applications (supplemented in Roth, 2002).

In the large majority of variants and applications of the assignment game, the goal has remained the same: to create two-sided matches that its members prefer over any other match they could achieve. In game terms, they sought to create matches that were in the core, which meant that no member of any matching pair could form a match with any other market participant that they preferred more.

²⁷ Their suggestion that indivisibilities in production and what they call "human existence" may have a role in "explaining such interesting realities as large and small cities" also bears further exploration.

3.3 *Game theory: Top trading cycle*

In their 1962 work, Gale and Shapley also briefly mentioned what they called “the problem of the roommates”. They outlined the problem of pairing up two elements drawn from a single set—in the prototypical case, a set of boys looking to pair up as roommates.

Shapley and Scarf (1974) expanded this idea, developing the one-sided market model. They posited a market consisting of a set of agents who each own a house and (for unspecified reasons) wish to trade their houses amongst themselves.²⁸ They reported that Gale had developed an algorithm, dubbed the “Top Trading Cycle” process (TTC), which allotted the houses amongst the traders, based on their preferences. TTC was a process to facilitate the trading of indivisible goods within a group, in its original form without the use of money. Like the assignment game, its goal was to produce stable matches.

As with the first two approaches, considerable work has been done in the past forty years to expand this model (Konishi *et al.*, 2001, provides a brief overview). However, compared to the other two approaches, relatively little work has been done to apply this model and its variants to real-world situations. (See Wang and Krishna, 2006, for a proposed application).

3.4 *Comparing approaches*

Work has also been done (including Quinzii, 1984; Quint, 1997; and Krishna and Wang, 2007) to explore the relationships between these three approaches, and to classify them using various criteria.

The intertwined history of these three approaches, and that fact that they are all designed to find optimal matches for indivisible goods, may suggest that the outcomes resulting from the three approaches would be similar, or perhaps even identical. The outcomes of the approaches are identical in particular instances (as both Von Neumann, 1953, and Shapley and Shubik, 1972, pointed out), but not in general.

3.4.1 *An example comparison*

Let the four elements of set $A = \{a_1, a_2, a_3, a_4\}$ each start with a corresponding indivisible good of the set $H = \{h_1, h_2, h_3, h_4\}$. Each element of A assigns a utility function to every element of H with a real number of the set $[0, 1)$:

²⁸ In the literature which followed, these indivisible goods began to be referred to as heterogeneous (*e.g.*: Crawford and Knoer, 1981, on heterogeneous firms and workers), and then as heterogeneous houses (*e.g.*, Gerber, 1986; Osborne, 2004). In part, this may have been to distinguish this problem from the Böhm-Bawerk (1891) market in ‘homogenous horses’. The term “heterogeneous houses” is not as clear as it might be. It is the utilities which are heterogeneous; the houses may or may not differ.


Table 2. Example utility functions.

	h_1	h_2	h_3	h_4
a_1	0.7	0.9	0.8	0.6
a_2	0.6	0.8	0.9	0.1
a_3	0.8	0.6	0.4	0.7
a_4	0.9	0.4	0.3	0.7

3.4.2 Assignment problem approach

Using these numbers in a classic assignment problem first requires some interpretation of the utility functions. Let the set A be workers, and the set H be the machines each is currently using to perform work. Their work is currently providing a utility to the firm²⁹— a_1 's work has a utility of 0.7, a_2 's work a utility of 0.8, a_3 's a utility of 0.4, and a_4 's a utility of 0.7, for a total utility to the firm of 2.6. However, each worker would also provide a utility if they used a different machine. The utility of the work of a_1 would be 0.9 if they used machine h_2 , 0.8 with h_3 , and 0.6 with h_4 . The utility of the each remaining worker/machine combinations is listed in the rows that follow. The question to be answered is: Should the workers exchange machines? The assignment problem approach would say that, indeed, each worker should be using a different machine— a_1 using h_2 , a_2 using h_3 , a_3 using h_4 , and a_4 using h_1 :

Figure 5. Assignment problem allocations.



3.4.3 Assignment game approach


In a two-sided market, we might say that a_1 and a_2 are members of one class (perhaps sellers), while a_3 and a_4 are of another class (perhaps buyers). Given the example's utility matrix, a_1 would be willing to trade their good for either h_2 or h_3 , but not for h_4 .³⁰

²⁹ "...providing...to the firm" is a convention in optimization literature. It does not imply that the entire utility remains with the firm and that none of it returns to the worker. How the utility is divided between the firm and the worker once it is created is a separate consideration.

³⁰ Even if the rules of the market enabled them to do so (permitting trading between sellers), a_2 would refuse to trade with a_1 , because they value h_1 less than their own initial good. As well, even if the market enabled the two buyers— a_3 and a_4 —to trade, they would not do so; a_4 values a_3 's good less than their own, while a_3 would prefer to trade with a_1 over a_4 .

After considering all options, a_1 would offer to trade with a_3 . After considering all options, a_3 would accept a_1 's offer, preferring h_1 over all others. After a_1 and a_3 trade and leave the market, a_2 and a_4 would not be motivated to trade with each other, and would each leave with their original good.

Figure 6. Assignment game allotments.




3.4.4 TTC approach

In the TTC approach, we take the preferences outlined in the assignment game approach and examine all possible trading cycles. We discover that a_1 can be part of a stable 3-cycle, because they prefer h_2 most, while a_2 prefers h_3 most, and a_3 prefers h_1 most.

Following the TTC procedure, the three trade and leave the market. Unable to trade, a_4 forms a 1-cycle and leaves with their original endowment:

Figure 7. Top Trading Cycle allotments.



The result is stable because no coalition of A can form a cycle which its member(s) would prefer over the arrangement the TTC protocol creates.

3.5 Differences

In this example, each approach produces a different outcome:

Table 3. Assignments resulting from each approach.

	a_1	a_2	a_3	a_4
assignment problem	h_2	h_3	h_4	h_1
assignment game	h_3	h_2	h_1	h_4
TTC	h_2	h_3	h_1	h_4

(Of course, the three different methods do not always produce different assignments.)

Because the assignments may be different between the three approaches, the benefits may also be different. In this particular example, assuming a costless market:

Table 4: Benefits of each approach.

	b_1	b_2	b_3	b_4	benefit of market round (B)
assignment problem	0.2	0.1	0.3	0.2	0.8
assignment game	0.1	0.0	0.4	0.0	0.5
TTC	0.2	0.1	0.4	0.0	0.7

The differences arise because, despite their similarities, the assumptions, goals, and rules of the three approaches are different.

3.5.1 Agency differences

In the two-sided matches of the assignment game, both participants are treated as self-interested decision-makers—they are both agents. Both partners in a match must be satisfied—they must benefit as much or more from the match than from any available alternative—if the assignment is to be stable.

In the one-sided matching of TTC (at least in its basic form), only one side of the match is a decisionmaking agent; the house is indifferent to who lives in it.

In the assignment problem, neither side in the assignment is a decision-maker; workers are not agents. While workers may prefer one machine over another, their preferences are not considered in making the allocations.³¹ In the assignment problem's classic form, stability is imposed by a (usually unmentioned) central decision-maker. From a game-theoretic perspective, the assignment problem models a single-person game.

³¹ In some variants, assignment tools been adapted to at least partially accommodate some agency by those being allocated, for example, in Caron *et al.* (1999) and Pápai (2000).

3.5.2 *Differing advantages*

A crucial advantage of the assignment problem is its ability to maximize benefits. Because the assignment problem maximizes for only one decision-maker, it is less constrained in its outcomes than the other two approaches. As a result, when benefits can be compared, it will always yield a benefit equal to or greater than the other two approaches.

Much discussion about markets in the economic literature is focused on the causes of (and possible solutions to) market failures. The assignment game reverses that perspective; providing a quantifiable measure of what market *success* looks like, and a rigorous method for achieving that success. A crucial advantage of the assignment game is its ability to model the most beneficial stable matches achievable in traditional, bilateral trading between agents.

A crucial advantage of the TTC is its release from the requirement to match traders into pairs; it is a revolutionary idea. As a result, all else equal, TTC will always produce benefits equal to or greater than the assignment game.

3.5.3 *Limitations of applicability*

While elements of all three approaches are useful in developing a model for an ongoing market in heterogeneous indivisible goods, none of them model such a market fully, nor do they model market or rule preferences and market-choice behaviour.

A limited amount of work (see Eveborn *et al.*, 2006) has been done on repeated assignments from an assignment problem perspective.³² It is not surprising that the one-shot assignment problem remains the norm; a single decision-maker makes an optimal allocation with each assignment. Unless there are time or capacity constraints,³³ considering assignments in sequence rather than separately probably will not increase utility.³⁴

³² Albright (1974), Albright and Derman (1972) Derman *et al.* (1972), Kennedy (1986), David and Yechiali (1995), Kaplan (1986), and Richter (1989) examine sequential assignments (where, unlike the classic assignment problem, not all assignments are made at a single point in time). However, sequential assignments are not repeated ones; no element ever becomes available for future assignments. (Although Kaplan does not use a game-theoretic model, the problem considered is similar to the ones considered by the assignment game literature. In Kaplan, housing is (sequentially) allocated to people. In contrast to the assignment game literature, the housing is considered to be homogenous, but the people being assigned are not; they are prioritized based on need. More significantly, the people given the housing are not agents; the housing is allocated by a central decision-maker.)

³³ Time and capacity constraints are addressed as bottleneck and queuing problems, which are outside the problem at hand.

³⁴ It may, however, be relevant if utility fluctuates with time—a worker may not perform all tasks in a sequence equally well, or equipment may wear and become less accurate during use. These considerations are beyond the scope of this paper.

What is modelled here is analogous to a repeated game while, except for Benabou (1993), Jackson and Watts (2005), and Lennon *et al.* (2007), the literature of assignment games and TTC focuses largely on non-repeated games.³⁵ This may be, in part, because “houses” are often used as the nominal goods in these games. As a result, a second extraneous quality—non-consumability—is bundled with the good’s essential quality of indivisibility. If a player needs only one good at a time, and if it cannot be consumed, the player will never need another, so the game ends once the allotments are made.

However, as the literature on repeated games makes clear (see Mailath and Samuelson, 2006), the incentives for agents in repeated games differ fundamentally from those in non-repeated ones, even if the rules governing each stage in a repeated game are identical to those of a non-repeated game.

4 *n*-trading formulation

The intention of this formulation is to integrate the merits of the three approaches—to develop trading rules which have the utility-maximization qualities of the assignment problem, which models agents’ self-interest, and which facilitate *n*-cycle trading—and to extend those merits onto an ongoing market in heterogeneous indivisible goods. It is derived from the maximization form of the assignment problem, with the addition of a trading cost, and a restraint that might be called a “motivation requirement”:

Maximize

$$\sum_{i=1}^n \sum_{j=1}^n x_{ij} (v_{ij} - t_i - v_{ii})$$

subject to:

$$\sum_{i=1}^n x_{ij} = 1$$

$$\sum_{j=1}^n x_{ji} = 1$$

$$\sum_{j=1}^n (x_{ij} (v_{ij} - t_i) - t_i x_{ii} - v_{ii}) \geq 0$$

$$x_{ij} \in \{0,1\}, \quad i = 1, \dots, n \\ j = 1, \dots, n \text{ where } n \text{ is the market size.}$$

$$v_{ij} \in [0, 1)$$

$$t_i \in [0, 1)$$

³⁵ Klaus and Miyagawa (2002) consider the situation where agents may desire to consume more than one good, but they model it as a non-repeating game. Llorca *et al.* (2003) and Llorca *et al.* (2004) deal with semi-infinite and infinite assignment games, but the infinities have to do with the number of agents, not the number of times agents are assigned.

The first two capacity constraints ensure that every agent ends each market round with one good.

The third constraint ensures every agent who trades is motivated to do so, because they will only trade if the utility of their allotted (new) good is at least equal to the utility of their old (endowed) good, plus the cost of trading.

5 A comparison of existing market models with n -trading

Let us call the class of all markets which facilitate only bilateral trading (whether using assignment game rules or not) *Traditional* markets (T -markets: m_T), and say they are operating under *Traditional* rules (T -rules: r_T). Because David Gale invented the top trading cycle, let us call the class of all markets using the top trading cycle mechanism to allot goods *Gale markets* (G -markets: m_G), and say they are using G -rules (r_G).³⁶ Let us call the class of all markets using the n -trading formulation given above N -markets.

While we cannot compare the outcomes and benefits of all t -, g -, and n -markets, we can compare the outcomes and benefits of a particular modelled market operated under t -, g -, or n -rulesets. The n -ruleset would be derived from the general n -market formulation given above, while the g -ruleset would be derived from Shapley and Scarf (1974). We need to specify a particular t -ruleset. To facilitate comparison, let us make the least possible change to the n -ruleset given, and specify only that all trades must be bilateral. This requires the addition of one constraint to the n -ruleset formulation.³⁷

$$x_{ij} = x_{ji}$$

5.1 Making the comparison

A computer simulation using Premium Solver was conducted to compare the benefits of 4-agent t -, g - and n -markets. 98 rounds were conducted, with four potential traders in each round, using randomly-generated values from the set $[0,1)$ for all utility functions. The trading cost for all agents was held constant ($t = 0.10$) throughout.

Of those 98 market rounds, 88 (90%) were viable (there were at least two motivated traders in a round).

A set of utility functions was randomly generated for each round. Each round used the step series of section 2.5.3, and was run three times—once with a t -ruleset, once with a g -ruleset and once with an n -ruleset—creating three parallel markets.

³⁶ Although the rules are consistent throughout the literature, there does not appear to be a standard term for these markets and the rules which govern. Writers refer to “Gale’s Top Trading Cycles Method” (Abraham, *et al.*, 2004), “house exchange markets” (Ehlers, 2002), “the n -player house-swapping game of Shapley and Scarf” (Quint and Wako, 2004), “the Shapley-Scarf...housing market” (Ma, 1994), and “the Shapley-Scarf economy” (Konishi, *et al.* 2001).

³⁷ The need for this additional constraint explains why one-sided markets are, in general, more beneficial than t -markets; two-sided markets are a more-constrained version of one-sided markets.

There were small but important differences in outcomes between the *t*- *g*- and *n*-markets in the simulation.

Table 5. Summary results for all potential traders under *t*- *g*- and *n*-rules in a 4-agent market.

	rules		
	<i>t</i>	<i>g</i>	<i>n</i>
successful rounds	67	68	70
successful traders	143	167	181
average utility of all endowed goods at start of market round	0.48		
average utility of all allotted goods at end of market round	0.64	0.66	0.67

Not all potential traders were motivated to trade. Only those who found a good on offer they judged to be significantly better than their own would have an interest in trading.

Table 6. Summary results for motivated traders under *t*- *g*- and *n*-rules in a 4-agent market.

	rules		
	<i>t</i>	<i>g</i>	<i>n</i>
% of motivated traders able to trade	54%	63%	68%
average benefit of trading to motivated traders	65%	72%	76%

These differences all follow from the differences in constraints—on average, the fewer the constraints, the greater the benefits.

If these agents were able to choose which of these three markets to participate in (*i.e.*: if they were engaged in a meta-game), they would prefer *t*-markets (and *t*-rules) less than the other two options.

In 17 of the rounds (17%) *n*-rules produced different allotments than *g*-rules. In each case, the total benefit of that round was greater under *n*-rules than *g*-rules. However, in each case, the outcome under *g*-rules was stable, while under *n*-rules it was unstable. This trade-off between benefit and stability means that, all else equal, impatient (high- δ) agents will be more likely to prefer *g*-rules, while patient (low- δ) agents will be more likely to prefer *n*-rules.³⁸


³⁸ Where the crossover point is (what the value of δ is at which *n*-rules and *g*-rules are equally beneficial) cannot be stated in general. It will depend on market- and rule-specific factors.

6 Computer simulation comparing n -trading and traditional trading

Computer simulations were conducted, comparing the benefits of n -markets and t -markets ranging in size from 2-agent markets to 35-agent markets.³⁹ Enough rounds were conducted so that every size market had 70 viable rounds. All other factors were the same as they had been in the computer simulation of section 5.2. (See Online Appendix 1 for computer simulation results. Results of all rounds for all markets available on request.)

6.1 Benefit of trading

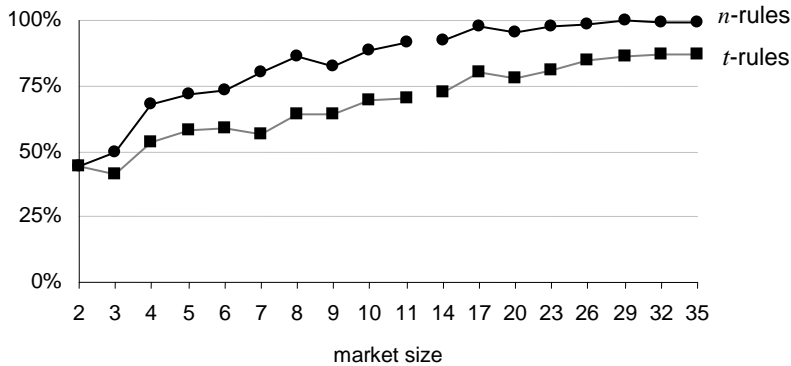
Figure 8. Average benefit.



On average, agents benefited more from n -rules than t -rules.

6.2 Ability to trade

Figure 9. Motivated traders able to trade.




More motivated traders were able to trade (become successful traders) under n -rules than under t -rules.

³⁹ Market sizes were the natural numbers from 2 to 11 (2, 3, 4, 5, 6, 7, 8, 9, 10, and 11) and, from there, every third natural number to 35 (14, 17, 20, 23, 26, 29, 32, and 35).

6.3 Getting most preferred good

Figure 10. Motivated traders who got their most-preferred good.

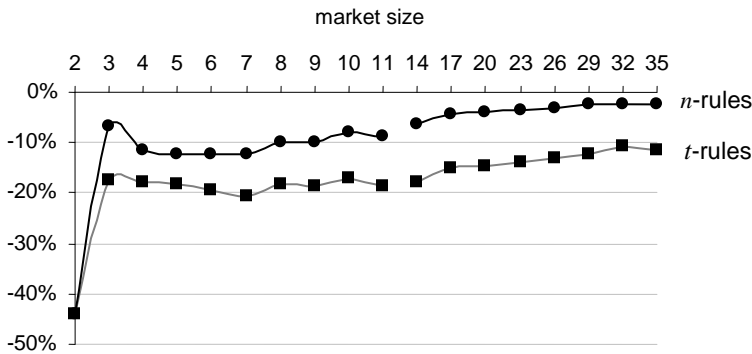


At least some motivated traders got their most-preferred good (first choice amongst goods on offer) in at least some rounds at all market sizes. However, this result was consistently higher under *n*-rules than under *t*-rules. The difference increased as the market size increased.

6.4 Comparison to an Ideal Optimum market

If all agents in a round end that round with their most-preferred good, that market round has delivered all agents their ideal outcome—an *ideal round*. The total benefit of an ideal round is the sum of the benefits of the goods most preferred by each participant; that total is that round's *Ideal Optimum* (IO). An *ideal market* would be one in which all participants ended every round with their most preferred good. Of course, no ideal markets exists, but both real markets and simulated markets can be judged against this ideal.

Figure 11. Shortfall from Ideal Optimum.

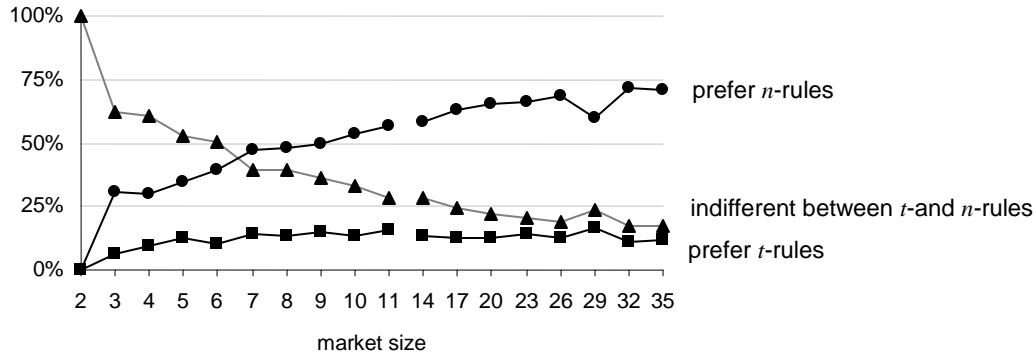


n-rules delivered benefits closer to the IO than *t*-rules. The ratio difference between the shortfalls increased as the market size increased.

6.5 Preferred rules

Depending on the utility functions of a particular round, some agents received a greater benefit under t -rules than under n -rules, in others the reverse was true, while in others they received the same benefits under both rules.

Figure 12. Rules preferred by agents.



If agents expect multiple market encounters, or if their preferences are based on the probability of benefit, they should prefer n -rules. This preference should hold no matter the market size, but should be stronger with larger markets.

6.6 Results

The results of the computer simulation strongly suggest that n -rules will be preferred by agents over traditional t -rules.

7 Testing practicality of n -trading in facilitated prototype markets

For rules to be useful in actual markets, it must be practical to apply them. Calculations must be completed quickly, the process must be able to accommodate a variety of situations, and participants must see their benefits. Three facilitated prototype markets, each consisting of a single round, were conducted to test whether n -trading is practical.

Rather than having utilities functions be randomly generated, substitutes for utility were produced by test participants. The resulting data was run through Premium Solver using variants of both n -rules and t -rules drawn from the computer simulations.⁴⁰

Eight participants who like (and occasionally collect) wine were brought together. Each participant brought three different bottles of wine to trade. All three markets included all eight participants.⁴¹

⁴⁰ See *Online Appendix 6* for a comparison between prototype and computer simulation markets. Because the prototype markets were extremely limited in scope and because the valuations generated were not utility functions, while interesting, this comparison is not statistically valid.

⁴¹ In addition to the three bottles intended for trade, each participant was asked to bring one bottle to sample with others; this bottle matched one of the wines they were intending to trade. As a result, all participants

7.1 Market 1 (rank ordering): Comparing t_1 - and n_1 -rules

In the first market, participants were directed to rank all eight wines available for trade (including their own) in their order of preference. No duplicate rankings were permitted, and they were not permitted to omit any available bottles from the rankings. The rules used in the computer simulations were adapted to maximize total rank movement rather than maximize benefit. The results (Online Appendix 2) indicate that n_1 -rules resulted in trading allotments that improved the rankings more than t_1 -rules did.

7.2 Market 2 (scrip offers): Comparing t_2 - and n_2 -rules

In the second market, participants were given 100 units of “play-money” scrip⁴² (\$100). Participants used the scrip to offer to “purchase” bottles from other participants.

In this market, an offer to purchase by a participant had two components—offering to relinquish their bottle, and offering no more than \$100. As a result, the maximum difference in scrip valuation they could set between their bottle and a bottle they wanted was \$100.

Participants were not required to value all bottles. Both sealed-bid valuations and open-bid valuations were solicited from all participants. The rules used in the computer simulation were adapted to maximize total scrip valuation increases for participants.

Although sealed-bid (Online Appendix 3) and open-bid (Online Appendix 4) scrip valuations were not identical, the results were generally similar. In both closed-bid and open-bid valuations, n_2 -rules resulted in trading allotments that improved the scrip valuations more than did t_2 -rules.

7.3 Market 3 (money offers): Comparing t_3 - and n_3 -rules

In the third market, instead of scrip, money (Canadian dollars) was used. Each participant brought \$10 of their own funds, and was told that they would be asked to pay half the difference between the value they assigned to their own bottle and the value of any bottle they received in trade. The linear program again maximized total valuation increases.

The n_3 -rules enabled participants to increase the value (to them) of the goods they owned by roughly 50%, while t_3 -rules enabled them to increase the value of their goods by roughly 20% (Online Appendix 5).

were able to sample 1/3 of the wines available for trade. Only one direction was given to participants as to what wines they should bring for trading: They were asked to bring a bottle they valued between \$10 and \$20 for the third trading market. Each market included an opportunity for participants to promote the wine they were offering for trade. Participants were free to discuss the wines and their valuations throughout.

⁴² Scrip is a medium of exchange useable within a specific market, but having no value outside that market.

7.4 Results

All three markets (including explanations of the process and rules of each market, “sales pitches” between participants on the merits of the wines on offer, participants’ recording of preferences, data entry, computer calculations, and the actual exchange of bottles of wine) were completed in two hours. The results suggest that facilitated n -trading is practical.

8 Potential applications

If the differences between n -trading and traditional bilateral trading are replicated in real-world markets, both participants and market organizers should prefer n -rules. As well, because the benefits generated are, on average, greater, if trades are monetized, the revenues generated in real-world n -markets should also be greater.

The most fruitful applications should be in markets where valuations of the goods involved differ markedly between participants.⁴³ Because n -trading is a new approach to markets, the higher the level of trust between participants, the more likely n -trading is to be adopted. Trust is likely to be high if either relationships that precede or supersede the market exist, or if participants expect multiple market encounters. In low-trust markets, participants may be more likely to prefer g -rules.

While n -trading will not be applicable to the most prominent form of trading—two-sided matches between a buyer with money and a seller with goods—real-world applications of the n -trading model may still be broad. Some possible applications of n -trading include trading in tangible goods (*e.g.*: between collectors), trading in less tangible goods (*e.g.*: time-share arrangements, airport terminal rights, perhaps electricity trading between power companies), and trading in services (*e.g.*: trading player contracts between sports franchises, trading shifts between workers). It may be particularly applicable to heterogenous resource and space use trading systems being proposed under ocean zoning regimes (Edwards, 2008).

Appendix

Commonalities

To facilitate comparison between markets and rules, the markets modelled in this paper, the agents acting in them, and the goods being traded in them, all have a number of features in common.

All these modelled markets:

(1) are *facilitated*

⁴³ If all participants agree on their valuation of all the goods on offer (in the view of the participants, goods are undifferentiated), n -trading will provide no more benefit than the traditional approach.

A central clearinghouse exists for goods in a market, to which participants disclose their valuations.

(2) are *confidential*

The clearinghouse will not disclose the valuations of a participant to any other participant.

(3) are *voluntary*

Market rules will not require any participant to act against their interests.

(4) are *leisurely*

All agents are able to consider all options before any agent may act.

All goods are:

(1) *indivisible*

Less than a whole unit cannot be owned, used, or traded.

(2) *private*

Owned by exactly one participant prior to, and after, each market round.

(3) *transferable*

The owner of a good prior to a market round may or may not be the owner of that good after that round.

All agents are:

(1) *knowledgeable*

Aware of all rules of all markets.

(2) *competent*

Able to calculate the benefit of any action and any rule in any market)

(3) *selfish*

Will always take an action if and only if it benefits them.

(4) *non-greedy*

Can make use of only one good after a market round, and so will not leave a market round with more than one good.

(5) *tight-lipped*

Will not disclose information to another agent unless it benefits them.

(6) *compliant*

Obeys the rules of the market in which they are acting.

(7) *consistent between goods*

If $v_{ij} < v_{ik}$ and $v_{ik} < v_{il}$, then $v_{ij} < v_{il}$.

(8) *consistent between markets*

v_{ij} is the same in all markets.

All utility functions for all goods are:

(1) *strictly heterogeneous*

No agent will have the an identical utility function for two goods in a single market round.

(2) *personal*

The utility function of a specific good in a specific market round may vary between agents.

References

- Abdulkadiroğlu, A., T. Sönmez. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* **66** (3) 689-701.
- Abraham, D. J., K. Cechlarova, D. F. Manlove, K. Mehlhorn. 2004. Pareto optimality in house allocation problems. *Proceedings of ISAAC 2004: the 15th Annual International Symposium on Algorithms and Computation, 20-22 December, 2004* (R. Fleischer, G. Trippen, eds.).
- Albright, C. 1974. Optimal sequential assignments with random arrival times. *Management Science* **21** (1) 60-67.
- Albright, C., C. Derman. 1972. Asymptotic optimal policies for the stochastic sequential assignment problem. *Management Science* **19** (1) 46-51.
- Bausch D. O., G. G. Brown, D. R. Hundley, S. H. Rapp, R. E. Rosenthal. 1991. Mobilizing marine corps officers. *Interfaces* **21** (4) 26-38.
- Benabou, R. 1993. Search market equilibrium, bilateral heterogeneity, and repeat purchases. *Journal of Economic Theory* **60** 140-158.
- Biggart N. W., T. D. Beamish. 2003. The economic sociology of conventions: habit, custom, practice, and routine. *Annual Review of Sociology* **29** 443-464.
- Böhm-Bawerk, E. von. 1891. *Positive Theory of Capital* (trans. W. Smart). G. E. Steckert, 1923.

- Caron, G., P. Hansen, B. Jaumard. 1999. The assignment problem with seniority and job priority constraints. *Operations Research* **47** (3) 449 - 453
- Castaldo, S. 2008. *Trust in Market Relationships*. Cheltenham, UK: Edward Elgar.
- Crawford, V. P., E. M. Knoer. 1981. Job matching with heterogeneous firms and workers. *Econometrica* **49** (2) 437-450.
- David, I., U. Yechiali. 1995. One-Attribute sequential assignment match processes in discrete time. *Operations Research* **43** (5) 879-884.
- Derman, C., G. J. Lieberman, S. M. Ross. 1972. A sequential stochastic assignment problem. *Management Science* **18** (7) 349-355.
- Edwards, S. 2008. Ocean zoning, first possession and Coasean contracts. *Marine Policy* **32** 46-54.
- Ehlers, L. 2002. Coalitional strategy-proof house allocation. *Journal of Economic Theory* **105** 298-317.
- Eveborn, P, P. Flisberg, M. Rönnqvist. 2006. Laps Care: an operational system for staff planning of home care. *European Journal of Operational Research* **171** (3) 962-976.
- Gale, D., L. S. Shapley. 1962. College admissions and the stability of marriage. *American Mathematical Monthly*. **69** 9-15.
- Gardner, R., E. Ostrom. 1991. Rules and games. *Public Choice* **70** 121-149.
- Gerber, R. I. 1986. Efficiency of markets for heterogeneous, indivisible housing. *Regional Science and Urban Economics*. **16** 407-419.
- Hasker, K., 2007. Social norms and choice: a weak folk theorem for repeated matching games. *International Journal of Game Theory* **36** 137-146.
- Hayek, F. 1948. *Individualism and Economic Order*. University of Chicago Press.
- Hitchcock, F. L. 1941. The distribution of a product from several sources to numerous localities. *Journal of Mathematical Physics*. **20** 224-230.
- Howard, N. 1971. *Paradoxes of Rationality: Games, Metagames, and Political Behavior*. Cambridge, MA: The MIT Press.
- Jackson, M. O., A. Watts. 2005. Social games: Matching and the play of finitely repeated games. *FEEM Working Paper 38.05* (also Caltech Working Paper 1212). Accessed August 20, 2008 from <http://ssrn.com/abstract=688350>.
- Kandori, M. 1992. Social norms and community enforcement. *Review of Economic Studies* **59** 63-80.
- Kaplan, E. H. 1986. Tenant assignment models. *Operations Research* **34** (6) 832-843.
- Klaus, B., E. Miyagawa. 2002. Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory* **30** 421–435.
- Konishi, H., T. Quint, J. Wako. 2001. On the Shapley–Scarf economy: the case of multiple types of indivisible goods. *Journal of Mathematical Economics*. **35** 1-15.
- Koopmans, T. C. 1947. Optimum utilization of the transportation system. *The Econometric Society Meeting* (Washington, D.C., September 6-18, 1947; D. H. Leavens, ed.) reprinted in *Supplement to Econometrica* **17** (1949) 136-46.
- Koopmans, T. C., M. Beckmann. 1957. Assignment problems and the location of economic activities. *Econometrica* **25** (1) 53-76.

- Krishna, A., Y. Wang. 2007. The relationship between top trading cycles mechanism and top trading cycles and chains mechanism. *Journal of Economic Theory* **132** 539 - 547.
- Kuhn, H. W. 1955. The Hungarian Method for the assignment problem. *Naval Research Logistics Quarterly* **2** 83-97, reprinted in *Naval Research Logistics* **52** (1) 7–21.
- Lennon, C. G., J. M. McGowan, K. Y. Lin. 2007. A game-theoretic model for repeated assignment problem between two selfish agents. *Journal of the Operational Research Society* advance online publication, 7 November 2007 doi:10.1057/palgrave.jors.2602518.
- Llorca, N., S. Tijsy, J. Timmerlyz. 2003. Semi-infinite assignment problems and related games. *Mathematical Methods of Operations Research* **57** 67-78.
- Llorca, N., J. Sánchez-Soriano, S. Tijsy, J. Timmerlyz. 2004. The core and related solution concepts for infinite assignment games. *Sociedad de Estadística e Investigación Operativa* **12** (2) 331-350.
- Mailath, G. J., L. Samuelson. 2006. *Repeated games and reputations: long-run relationships*. New York, NY: Oxford University Press.
- Misztal, B. 1996. *Trust in Modern Societies: The Search for the Bases of Social Order*. Cambridge, UK: Polity Press.
- Miyagawa, E. 2002. Strategy-proofness and the core in house allocation problems. *Games and Economic Behavior* **38** 347–361.
- Niederle, M., A. E. Roth, T. Sönmez. 2007. Matching. *The New Palgrave Dictionary of Economics*, 2nd ed. Palgrave Macmillan.
- Öncan, T. 2007. A survey of the generalized assignment problem and its applications. *INFOR* **45** (3) 123-141.
- Orden, A. 1956. The transshipment problem. *Management Science* **2** (3) 276-285.
- Osborne, M. J. 2004. *An Introduction to Game Theory*. New York: Oxford University Press.
- Ostrom, E., R. Gardner, J. Walker. 1994. *Rules, games, and common-pool resources*. Ann Arbor, MI: University of Michigan Press.
- Pápai, S. 2000. Strategyproof assignment by hierarchical exchange. *Econometrica* **68** (6) 1403-1433.
- Pápai, S. 2003. Strategyproof exchange of indivisible goods. *Journal of Mathematical Economics* **39** 931–959.
- Plato. ca. 360 BCE. *Republic* (trans. B. Jowett). Accessed August 25, 2008 from <http://classics.mit.edu/Plato/republic.html>.
- Quint, T. 1997. Restricted houseswapping games. *Journal of Mathematical Economics* **27** 451-470.
- Quint, T., J. Wako. 2004. On houseswapping, the strict core, segmentation, and linear programming. *Mathematics of Operations Research* **29** (4) 861-877.
- Quinzii, M. 1984. Core and competitive equilibria with indivisibilities. *International Journal of Game Theory* **13** (1) 41-60.
- Righter, R. 1989. A resource allocation problem in a random environment. *Operations Research* **37** (2) 329-338.
- Roth, A. E. 2002. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* **70** (4) 1341–1378.

- Roth, A. E., M. A. O. Sotomayor. 1990. *Two-sided matching: A study in game-theoretic modelling and analysis* (Econometric Society Monographs, 18). Cambridge University Press, paperback edition 1992.
- Shapley, L. S. 1953. A value for n-person games. *Contributions to the Theory of Games II* (Annals of Mathematics Studies, 28). H. W. Kuhn, N. W. Tucker, eds. Princeton, NJ: Princeton University Press.
- Shapley, L. S. 1955. *Markets as Cooperative Games*. Santa Monica, CA: The Rand Corporation, P-629.
- Shapley, L. S., H. Scarf. 1974. On cores and indivisibility. *Journal of Mathematical Economics* **1** 23-37.
- Shapley, L. S., M. Shubik. 1969. On market games. *Journal of Economic Theory* **1** (1) 9-25.
- Shapley, L. S., M. Shubik. 1972. The assignment game I: The core. *International Journal of Game Theory* **1** 111-130.
- Smith, A. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations* (Edwin Cannan, ed.) London: Methuen and Co, 5th ed. 1904.
- Smith, Jr., L. W. 1955. Current status of the industrial use of linear programming. *Management Science*. **2** (2) 156-158.
- Svensson, L-G. 1999. Strategy-proof allocation of indivisible goods. *Social Choice Welfare* **16** 557-567.
- The Bretton Woods Agreements. 1944. *Decade of American Foreign Policy: Basic Documents, 1941-49*. Washington, DC: Government Printing Office, printed 1950.
- Von Neumann, J. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, vol. 2. (H. W. Kuhn, A. W. Tucker, eds.) Princeton University Press.
- Von Neumann, J., O. Morgenstern. 1944. *Theory of Games and Economic Behavior* (3rd ed., 1953). Princeton: Princeton University Press.
- Votaw, D. F., A. Orden. 1952. The personnel assignment problem. *Report on a Symposium on Linear Inequalities and Programming*. (A. Orden, L. Goldstein, eds.) 155-163.
- Wang, Y., A. Krishna. 2006. Timeshare exchange mechanisms. *Management Science*. **52** (8) 1223-1237.
- Wright, G. H. von. 1963. *Norm and Action: A logical enquiry*. London: Routledge & Kegan Paul.
- Yu, Y., V. K. Prasanna. 2003. Resource allocation for independent real-time tasks in heterogeneous systems for energy minimization. *Journal of Information Science and Engineering* **19** (3) 433-449.