Algorithmic Game Theory: Artificial Agents Play and Learn, Producing Surprising and Instructive Results

Steven O. Kimbrough
kimbrough (à) wharton.upenn.edu
http://opim-sky.wharton.upenn.edu/~sok/

23 January 2009

Carnegie Mellon University

Case Examples

Acknowledgements

I am most grateful for the contributions of many collaborators, especially Robert L. Axtell, James D. Laing, Fred Murphy, and David Harlan Wood. Of course, I claim exclusive rights to all errors contained herein.

Greetings and Goals

This talk: about simple computational models that bear on fundamental research into rationality and decision making in strategic contexts, and that even have prospects of being useful for applications.

Theme: simple learning in strategic contexts. Highlights from a research program.

First, I'll frame things a bit and set the context, including philosophical remarks about modeling. Then I'll present highlight results from a stream of research.

Structure of the Presentation

Part A:
Context & Background Briefly

Two Motivating Grand Challenges

- Theoretical challenge

  Foundations of game theory, social science generally: a successor to Rational Choice Theory.

- Applications challenge

  Social design: improve the scientific basis for design and management of social institutions, e.g., markets, education, . . .

From Binmore [Binmore, 1992, pages 50–1]

Here's what got me started.

> What is important here is that game theory does not pretend to tell you how to make judgments about the shortcomings [in terms of ideal rationality] of an opponent. **In making such judgments, you would be better advised to consult a psychologist than a game theorist.** Game theory is about what players will do when it is understood that both are rational in some [specific, very strong] sense. . . .

Context: Ways of studying strategic decision making

- A priori. Classical game theory. Game theory as a branch of applied mathematics. Now also "algorithmic game theory" in the sense of finding computational means to compute difficult equilibria.
- In vitro. Behavioral game theory. Experiments in the lab. Empirical study of play.
- In vivo. "Games in the wild." Observe and record.
- In silico. Evolutionary game theory. Epistemic game theory; computational heuristics. (i) Properties and behavior of strategies? (ii) of learning regimes? (Note: learning as randomized search. Fictitious play is not learning on this definition.)

From an interview with Robert Altman

Questioner:

Why do you hate Hollywood?

Altman:

I don't hate Hollywood. They make shoes. I make gloves. We're in different businesses.

Part B:
A Little Philosophy of Modeling

Classifying models: Two dimensions

- A dimension
  - Kinds of sources of a model's behavior
  - Four categories
- B dimension
  - Characterizing the model-world fit
  - Two categories

Best seen with examples. Note: impossibility of perfect distinctions.

Classifying Models: Dimension A1

- Category A1 $\approx$ associational; causal or fundamental

  Example: Newton's law of universal gravitation:

  $$F = G\frac{m_1 m_2}{r^2}$$

  Causal: the masses are thought to cause the force, they are not merely correlated or associated with it. Behavior in any particular case is a function of universal, measured factors, gravitation and G.

  Note: Independent from truth.

Classifying Models: Dimension A2

- Category A2 $\approx$ associational; not causal or fundamental

  Example: Hooke's law of elasticity:

$$\mathbf{F} = -k\mathbf{x}$$

  Not causal: merely associational, not universal. k differs by spring. Causation is involved, but no universality. Doesn't tell us *why* there is a force, merely measures it. "An engineering law, not a fundamental scientific law."

  Many other examples in social science, including: econometric models, regression models, utility/decision analysis models.

Classifying Models: Dimension A3

- Category A3 $\approx$ axiomatic

  Method (like Euclidean geometry): lay down axioms, assumptions, and derive conclusions.

  Examples: microeconomics. Cournot model.

  Not—needn't be—causal in the way that Hooke's law isn't either, maybe less so. (However, an argument for being implicitly causal.) Deduction as consistency checking.

  Many other examples, especially in microeconomics.

Classifying Models: Dimension A4

- Category A4 $\approx$ generative, constructive

  “If you can’t make one, then you don’t know how it works.”

  Compare: constructive versus non-constructive proofs.

  Examples: agent-based models; Darwin’s (and the modern) theory of evolution; cellular automata models.

  Typically computational; may be analytic.

Summary on Dimension A

- Four sources, or drivers, for model behavior (prediction, etc.):
  1. Fundamental associations among variables. (Gravity law).
  2. Non-fundamental associations among variables. (Hook's law; econometric models).
  3. Axiomatic setup. (Many microeconomic models; classical game theory).
  4. Procedures that generate behavior. (Agent-based models; evolution; etc.).

  Reminder: imperfect distinctions.

  Comment: "Explanation has to stop somewhere." A4 models are most satisfying modes of explanation.

Classifying Models: Dimension B

- Two categories for B:

1. B1 $\approx$ calibrated; parameters mapped to real data
   Examples: Newton's law of gravity, G; Hook's law, k
2. B2 $\approx$ uncalibrated; parameters mapped to real data
   Examples: Cournot model; Schelling's segregation model; basic Ising models; SOK's models, above (and below)

- Uncalibrated versus uncalibratable.

How Can Uncalibrated Models Possibly Be Scientifically Useful?

- Easy for Newton's law of gravitation or Hook's law of elasticity. Replace the calibrated parameter with a positive constant of proportionality.

$$\mathbf{F} \propto -\mathbf{x} \tag{1}$$

$$F \propto \frac{m_1 m_2}{r^2} \tag{2}$$

- Cournot (microeconomics broadly) teach "insight" with uncalibrated models.
- Similarly for Schelling's segregation model, and other A4 models.

Uncalibrated models may become calibrated.

Cournot Reference Model

Roughly:

- A market for a particular product supplied by n firms.
- During each time step each of the supplying firms offers quantity Q_i ($i = 1, 2, \ldots, n$) to the market, so that the total supply in a given period is

$$Q = \sum_{i=1}^{n} Q_i \tag{3}$$

- The unit price resulting is determined by the demand function—

$$P = \max\{a - \textit{slope} \times Q, 0\} \tag{4}$$

- Each firm i receives revenues of $P \times Q_i$.
- Firms may independently and without communication with each other adjust the quantities they offer to the market, their Q_is.
- In setting their Q_is each firm takes into account its unit cost of production, k_i, and the behavior of the other firms.
- Each firm follows the *best response* strategy
- If all of the firms do this they will reach the Cournot equilibrium in which the individual firm Cournot quantities are

$$Q_i^C(n, k_i) = \frac{(a - k_i)}{(n+1) \cdot slope} \tag{5}$$

This conclusion follows, via details not described here, mathematically.

“Insight”?

- Data $\rightsquigarrow$ patterns

  Note: “stylized facts”

- Models: mimic patterns

  Count as insight-productive if the imitation is acceptably accurate.

- Finding: uncalibrated models can be insight-productive.

- Point strengthened in the case of A4 models. Explanation. Development.

Part C:
Cases & Examples

Spatialized Prisoner's Dilemma, play C or play D

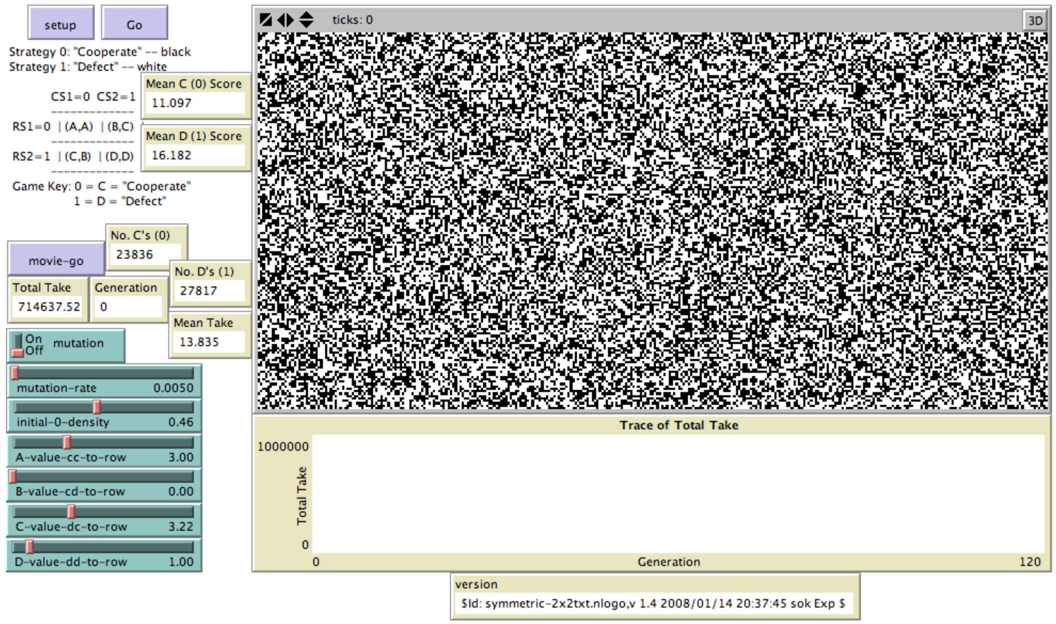

Figure 1: Prisoner's Dilemma initialized on the gridscape network.

Play your 8 neighbors; imitate the strategy of the most successful.

Cooperators devastated

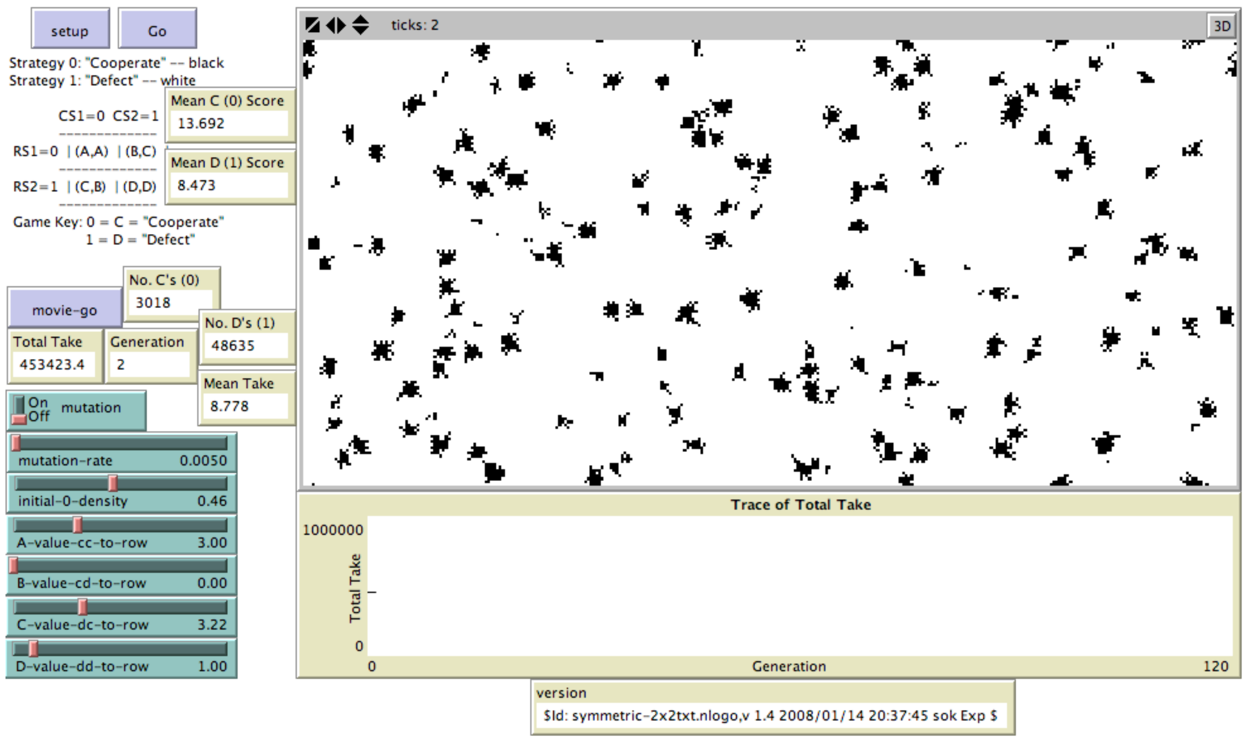

Figure 2: Prisoner's Dilemma generation 2: cooperators are decimated.

Surprise, surprise

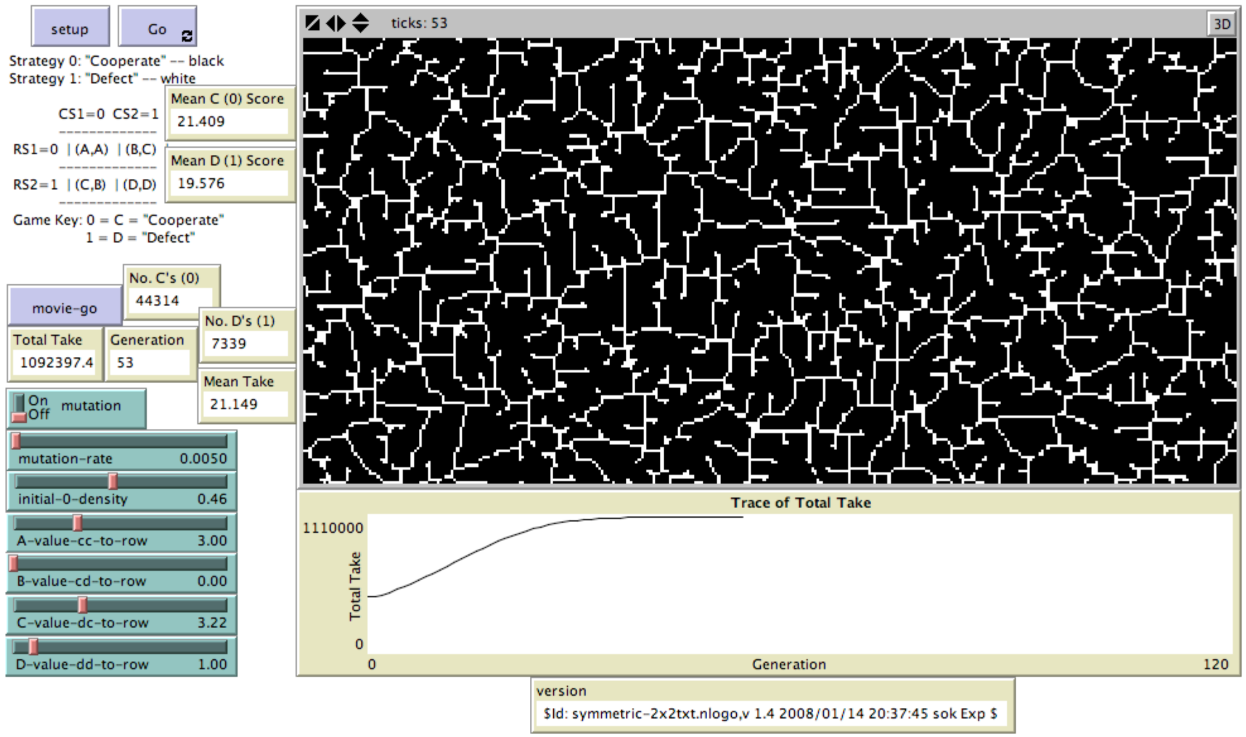

Figure 3: Prisoner's Dilemma generation 53: cooperators conquer.

And they stay that way

SOK's "invasion inequalities" for 2×2 games: In d dimensions, a hypercube of S_1s can invade an abutting field of S_2s if

$$A[2 \cdot 3^{d-1} - 1] + B3^{d-1} > \max\{D[2 \cdot 3^{d-1} - 1] + C3^{d-1}, D[3^d - 1]\} \quad (6)$$

Similarly, a hypercube of S_2s can invade an abutting field of S_1s if

$$D[2 \cdot 3^{d-1} - 1] + C3^{d-1} > \max\{A[2 \cdot 3^{d-1} - 1] + B3^{d-1}, A[3^d - 1]\} \quad (7)$$

	S_1		S_2	
S_1	A	A	B	C
S_2	C	B	D	D

A = 3, B = 0, C = 3.22, D = 1, d = 2

in our example, and $15 = 5A > \max\{5D + 3C, 8D\} = 14.66$.

Not a fluke. Important points.

- Algorithmic study of strategic interaction can yield surprising (and valid) results.
- Association between/among agents conduces cooperation. I call this the *shadow of society* (after Axelrod's *shadow of the future*)
- The regime of play matters greatly.

See [Kimbrough, 2008] for a brief presentation of the invasion inequalities and the Prisoner's Dilemma example.

Q-learning in 2×2 games

"Simple Reinforcement Learning Agents: Pareto Beats Nash in an Algorithmic Game Theory Study" [Kimbrough and Lu, 2005] http://opim-sky.wharton.upenn.edu/~sok/sokpapers/2005/kimbrough-lu-pareto-beats-nash.pdf

SSR: strategy selection regime. WE, wealth extracted is the proper measure of performance (ratio scale).

Backed off from Ideal Rationality and from Ideal Knowledge. We favor: LPS: learning in policy space.

Q-learning study. Finding: Pareto predicts better than Nash.

1. Alternative (or consideration) set.

   In the 2×2 case, conditioning on the last play by the counter-player, $\mathcal{A} = \{[0,0],[0,1],[1,0],[1,1]\}$ for each player. Interpretation of $[x,y]$: if on the last play the counter-player played x, on this play I play y, where 1=C and 0=D.

2. Attractiveness estimation.

   E.g., linear updating rule for $A^i, \ \ i \in \mathcal{A}$:
   $A^i_{t+1} = A^i_t + \alpha\{r^i_t - A^i_t\}$

   NewEstimate = CurrentEstimate + StepSize{reward - CurrentEstimate}

   In the case of Q-learning, $Q'(s,a) = Q(s,a) + \alpha(r - Q(s,a))$

3. Choice/exploration policy.

   ϵ–greedy. By default $\epsilon = 0.05$.

Figure 4: Summary of Elements in a Q-Learning Regime for 2×2 Games

Consider now play by two Q-learning agents of Indefinitely Repeated Prisoner's Dilemma. The stage game is repeated 100,000 times in each run. This counts as indefinite repetition because the agents have no, and can represent no, information concerning the length of play. We parameterize the stage game, as in Figure 5.

	C	D
C	(3,3)*	(0, 3+δ)*
D	(3+δ, 0)*	(δ, δ)#

Figure 5: A Parameterized PD Stage Game. #=Nash; *=Pareto

CC	CD	DC	DD	δ	Row's % CC
9422	218	183	177	0.05	0.963
9036	399	388	150	0.5	0.963
5691	738	678	2693	1	0.931
3506	179	275	6040	1.25	0.972
1181	184	116	8519	1.5	0.930
2	98	103	9797	1.75	0.805
97	114	91	9698	2	0.735
0	100	92	9808	2.5	0.839
2	96	94	9808	2.95	0.986

Table 1: Summary of Results for Prisoner's Dilemma. ε-greedy action selection. Totals are for the last 100 rounds of 100 series of 10,000 plays.

Cournot competition

- See “Learning to Collude Tacitly on Production Levels by Oligopolistic Agents” [Kimbrough and Murphy, 2009].

- Classic model of oligopoly. Here, firms offer quantities of a good and the market sets the price.

- The classic theory is undermotivated mathematically. Assumes “best response” behavior. This leads to the Cournot equilibrium, which lies between the monopoly quantity and the competitive quantity.

Cournot Reference Model

Roughly:

- A market for a particular product supplied by n firms.
- During each time step each of the supplying firms offers quantity Q_i ($i = 1, 2, \ldots, n$) to the market, so that the total supply in a given period is

$$Q = \sum_{i=1}^{n} Q_i \tag{8}$$

- The unit price resulting is determined by the demand function—

$$P = \max\{a - \textit{slope} \times Q, 0\} \tag{9}$$

- Each firm i receives revenues of $P \times Q_i$.
- Firms may independently and without communication with each other adjust the quantities they offer to the market, their Q_is.
- In setting their Q_is each firm takes into account its unit cost of production, k_i, and the behavior of the other firms.
- Each firm follows the *best response* strategy
- If all of the firms do this they will reach the Cournot equilibrium in which the individual firm Cournot quantities are

$$Q_i^C(n, k_i) = \frac{(a - k_i)}{(n+1) \cdot slope} \tag{10}$$

This conclusion follows, via details not described here, mathematically.

Is there another way?

- See [Kimbrough and Murphy, 2009], "Learning to Collude Tacitly on Production Levels by Oligopolistic Agents".

- Let's look at `http://opim-sky.wharton.upenn.edu/~sok/age/nlogo/oligopolyPutQuantity.html`.

Bertrand competition

- Now the firms compete on price.
- See "On Learning to Collude Tacitly on Price Levels by Oligopolist Agents" by Kimbrough and Murphy (`Bertrand-price-bidding.tex`).
- Economics theory: collusion is impossible. Even with just two firms in the market they will compete away their profits.

> If firm 1 really believes that firm 2 will charge a price $\hat{p}$ that is greater than the marginal cost, it will always pay firm 1 to cut its price to $\hat{p} - \varepsilon$. But firm 2 can reason the same way! Thus any price higher than marginal cost cannot be an equilibrium; the only equilibrium is the competitive equilibrium. [Varian, 2003, page 488]

Note the business literature on this: Don't do it!

- See `http://opim-sky.wharton.upenn.edu/~sok/age/nlogo/oligopolyBidPrice.html`

Management Implications?

- Executive pay? How to incentivize managers?

  By comparison with competitors? By returns to the firm?

- Once again, standard theory comes up short.

Electricity Prices (Bidding Supply Curves)

"Strategic Bidding of Supply Curves: An Agent-Based Approach to Exploring Supply Curve Equilibria" by Steven O. Kimbrough and Frederic H. Murphy. `http://opim-sky.wharton.upenn.edu/~sok/sokpapers/2009/Supplycurve_theorems_simulations.pdf`

- Firms bid a step function. The ISO (independent system operator) sums the bids horizontally and the market realizes the price. Not tractable analytically.
- Whereas extant mainstream literature tends to minimize the possibility of tacit collusion on prices, we find collusion all too easily established.

Electricity Auction

- Suppliers bid a step-function supply curve
- ISOs select steps until supply meets demand
- Everyone gets the market-clearing price
- Adjusted for location (which we ignore)

Models Used (in practice and the literature)

- Nash/Cournot
- Bertrand
- Supply-function equilibrium

The first two are used because they are tractable.

However, the third is what happens.

Problems with Analytical Approaches to Supply Function Equilibria

- Few theoretical results with most relying on overly simplistic assumptions

  Linear or continuous supply functions

  Restrictive boundary assumptions

- Multiple equilibria—which ones obtain?

- Repeated game with opportunities for learning and tacit collusion

Agent-Based Models

- Use agents that learn about their environment
- Agents can be policies in a games with many agents. Here the agents die or survive and reproduce
- Agents can choose actions/policies and select the actions/policies that improve their utility. Our approach.

Probe and Adjust

- Each agent has a cost curve
- In an episode each agent has a supply curve that it randomly adjusts to gauge the market consequences (sampling)
- After a fixed number of episodes (epoch), each agent adjusts the steps (prices and quantities) of the supply curve in the direction of improved outcomes

The agent does local search and moves in the direction of improved objective function

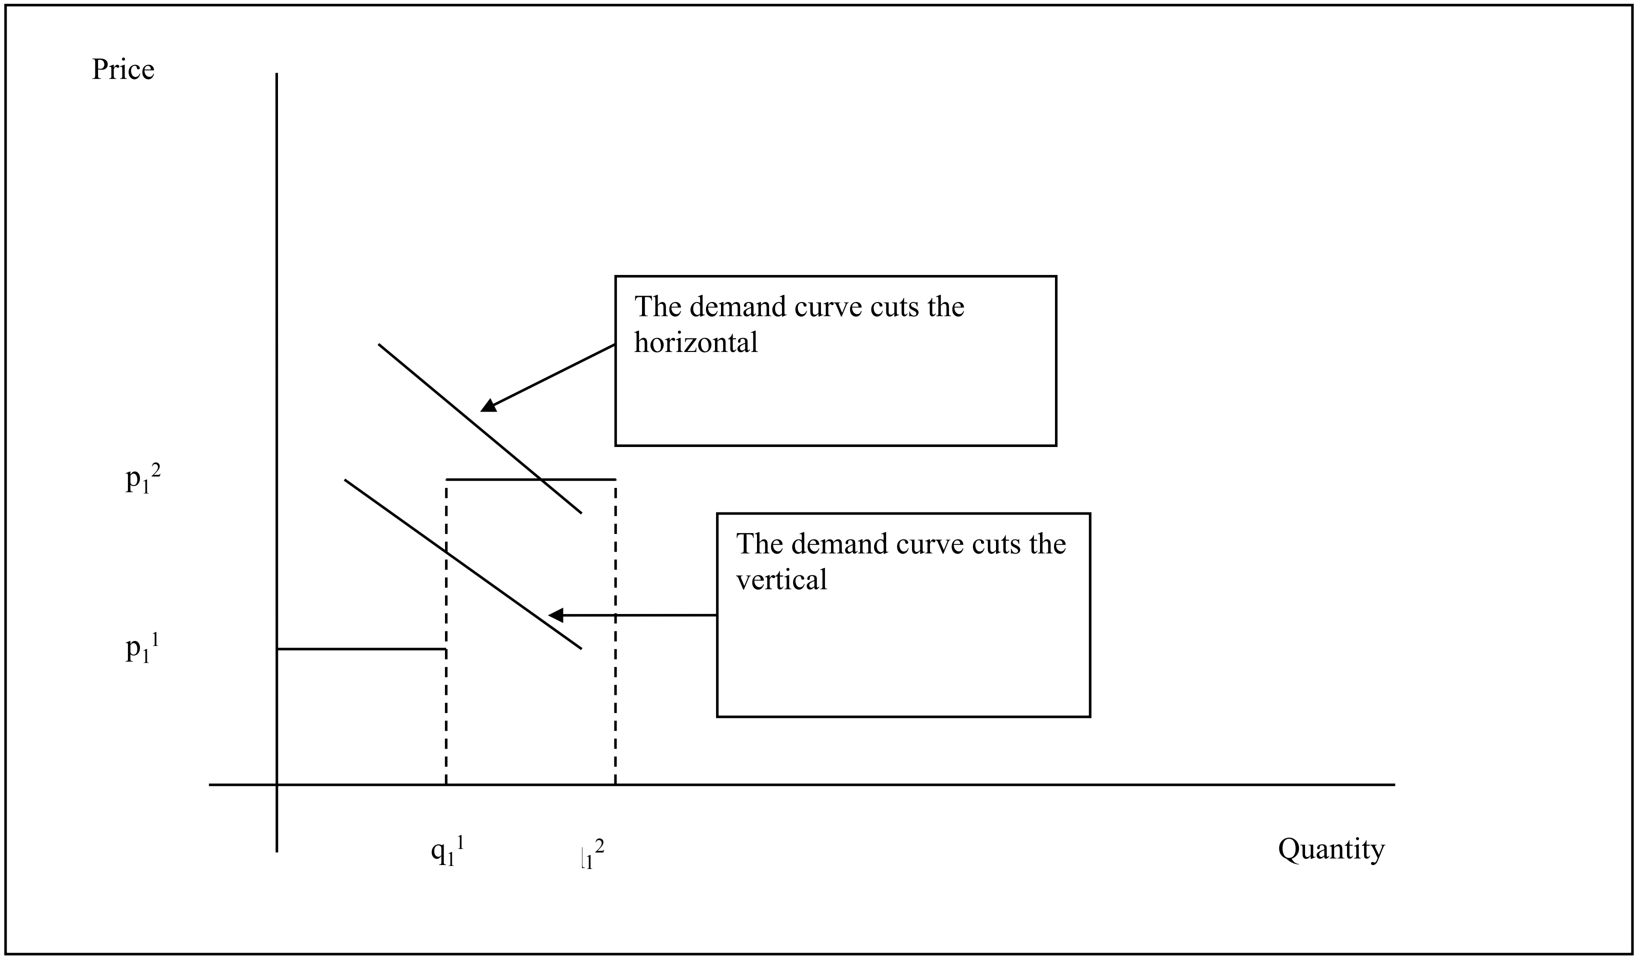

Figure 6: Player 1 definitions. Note that we have left out the supply from player 2, which would be included in the full market representation.

Theoretical Properties when both Players Maximize Firm Profits

- If both players are on verticals of offer curves, then both get to the Cournot equilibrium
- If both players are on horizontals of offer curves and their costs are different, then one eventually reaches a vertical

Why being "on a vertical" is advantageous to the suppliers.

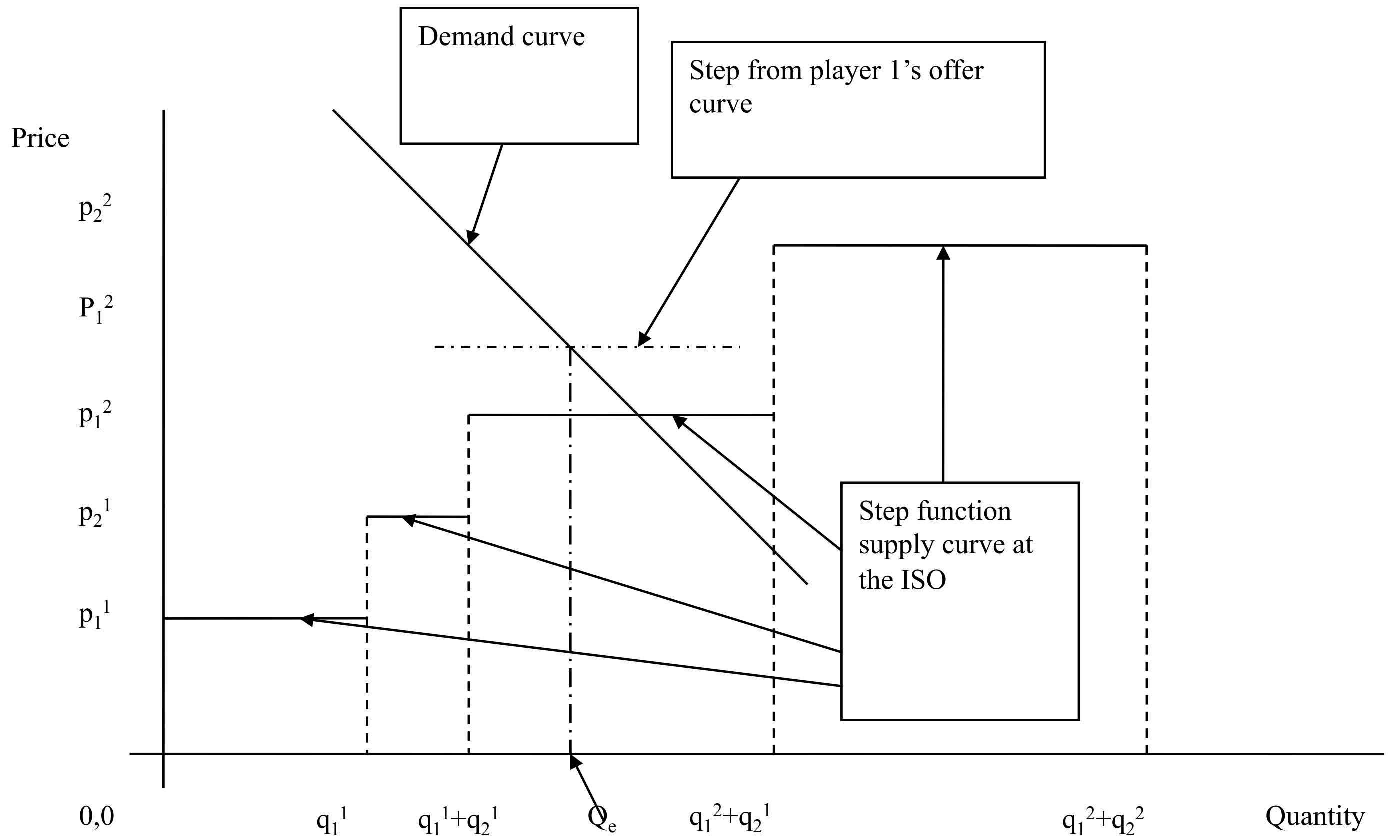

Figure 7: An equilibrium where the first steps of both players are utilized and player 1 raises its offer price to satisfy its Cournot equilibrium condition and the price P12 is below p22. [Small labeling errors.]

	(BT N, UT Own)	(BT N+1, UT Own)	(BT N, UT MRCOR)	(BT N+1, UT MRCOR)
(BT N, UT Own)	[38860] (24450, 14400) (21110, 17760)	[38970] (14520, 24120) (19150, 19820)	[38860] (14400, 24450) (18760, 20100)	[38920] (14510, 24140) (18490, 20430)
(BT N+1, UT Own)	[[(19820 , 19150)]]	[44350] (22190, 22140) (22190, 22160)	[39000] (24500, 14510) (24490, 14510)	[44930] (23550, 21310) (23570, 21360)
(BT N, UT MRCOR)	[[(20100, 18760)]]	[[(14510, 24490)]]	[48960] (24450, 24510) (24450, 24510)	[43100] (20530, 23230) (19750, 23350)
(BT N+1, UT MRCOR)	[[(20430, 18490)]]	[[(21360, 23570)]]	[[(23350, 19750)]]	[48950] (24050, 24900) (24320, 24630)

Table 2: RESULTS SUMMARY: Principal combinations of strategic positions, with sections of primary discussion indicated. 30 replications using the system clock to seed the random number generator. Row 1 of each cell is [mean `meanRunningIndustryProfits`]. Row 2 of each cell is (median `runningAvgProfitsFirm0`, median `runningAvgProfitsFirm1`). Row 3 of each cell is (mean `runningAvgProfitsFirm0`, mean `runningAvgProfitsFirm1`). When a row 3 entry is mark by [[. . .]] the entry is estimated by symmetry.

Agent-Based Model Built in NetLogo

- See `http://opim-sky.wharton.upenn.edu/~sok/age/nlogo/BidSupplyCurve.nlogo` or `.html`.

Conclusions on Supply Curve Study

- Our agents match classical behaviors in the classic models
- Without communicating but with an objective function that includes the community of firms they can tacitly collude
- The simulations are robust
- With agents it is possible to include more complicated features of markets like grids

Briefly: Two-sided matching

G-S, “triumph of game theory”, widely used

[Axtell and Kimbrough, 2008]

stability, fairness, social welfare

Agents can do better than G-S.

Anomalies for G-S

Several.

Key for us are three criteria by which to evaluate a matching:

1. Stability (G-S et al., [Roth and Sotomayor, 1990])

2. Social welfare (NB linear programming formulation)

3. Fairness (e.g., [Klaus and Klijn, 2006], [Fuku et al., 2006])

Stable matchings under G-S do not maximize social welfare and do not lead to fair outcomes.

Agent model and results

The challenge:

How might we find matchings that either (a) a stable and are improved wrt social welfare and/or fairness, or (b) are not stable, but not unacceptably so, and are improved wrt social welfare and/or fairness?

How, in particular, might we build agent-based models to investigate the problems of two-sided matching?

NB: Real-world complications often vex the deferred-acceptance approach.

Read the paper, but briefly. . .

Agent Model: Distributed Matching

- Two types of agents.
- Each individual of each type randomly (uniformly) ranks each individual of the other type.
- Each step a small number of agents, randomly chosen, become active and make proposals of engagement that are accepted or not.
- Steps accumulate into periods, when the number of agents activated equals the number of agents.
- With some probability, an engaged pair get married.

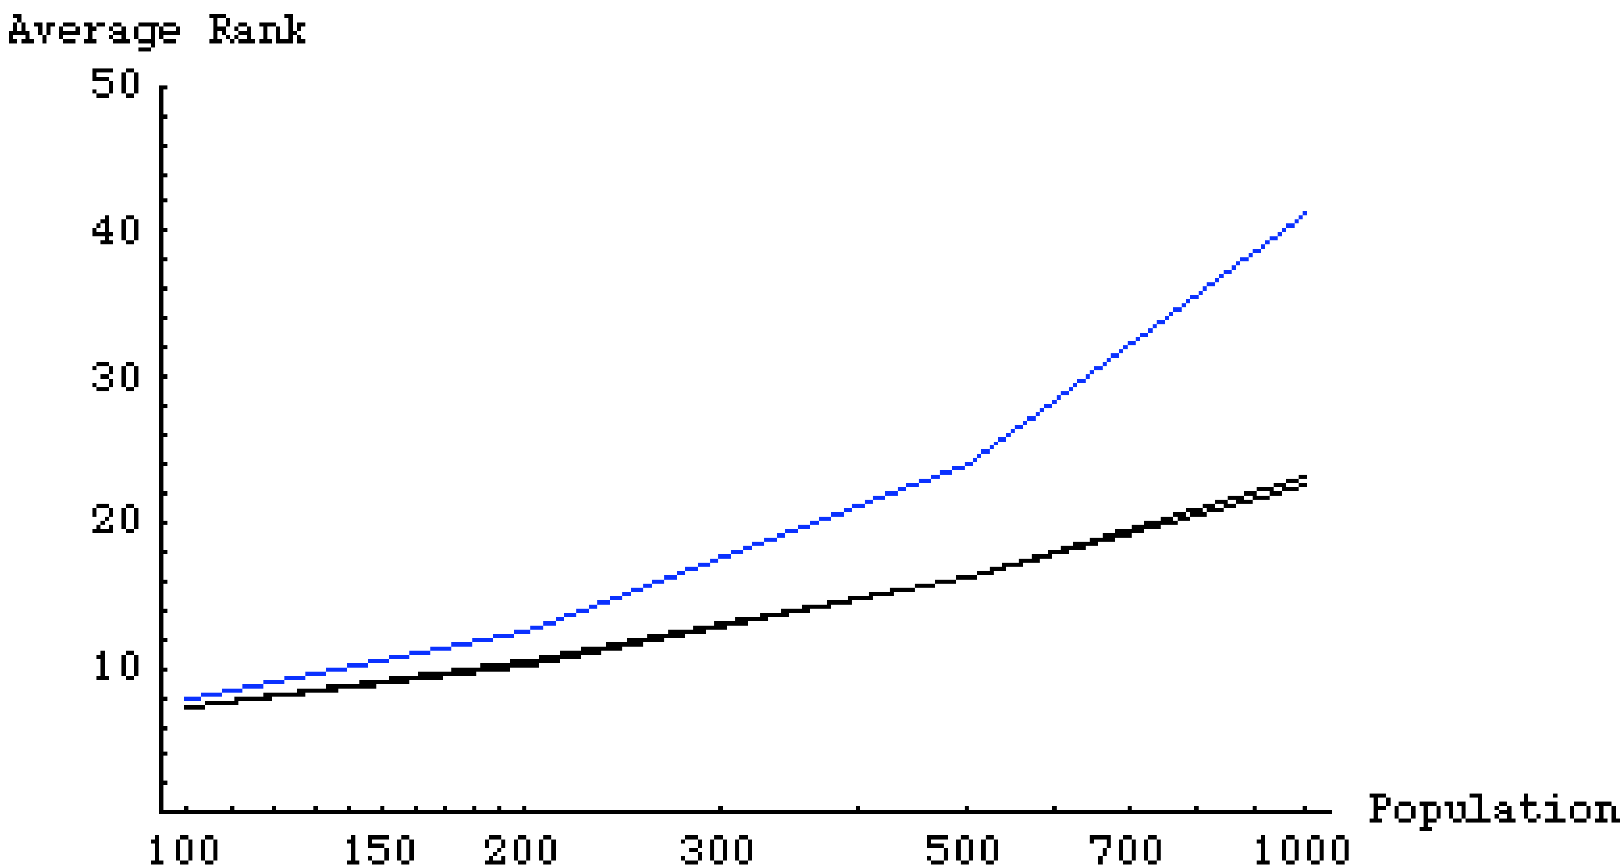

Figure 8: Figure 4 from paper. Compared to G-S (blue), our agents are better on SW and on fairness.

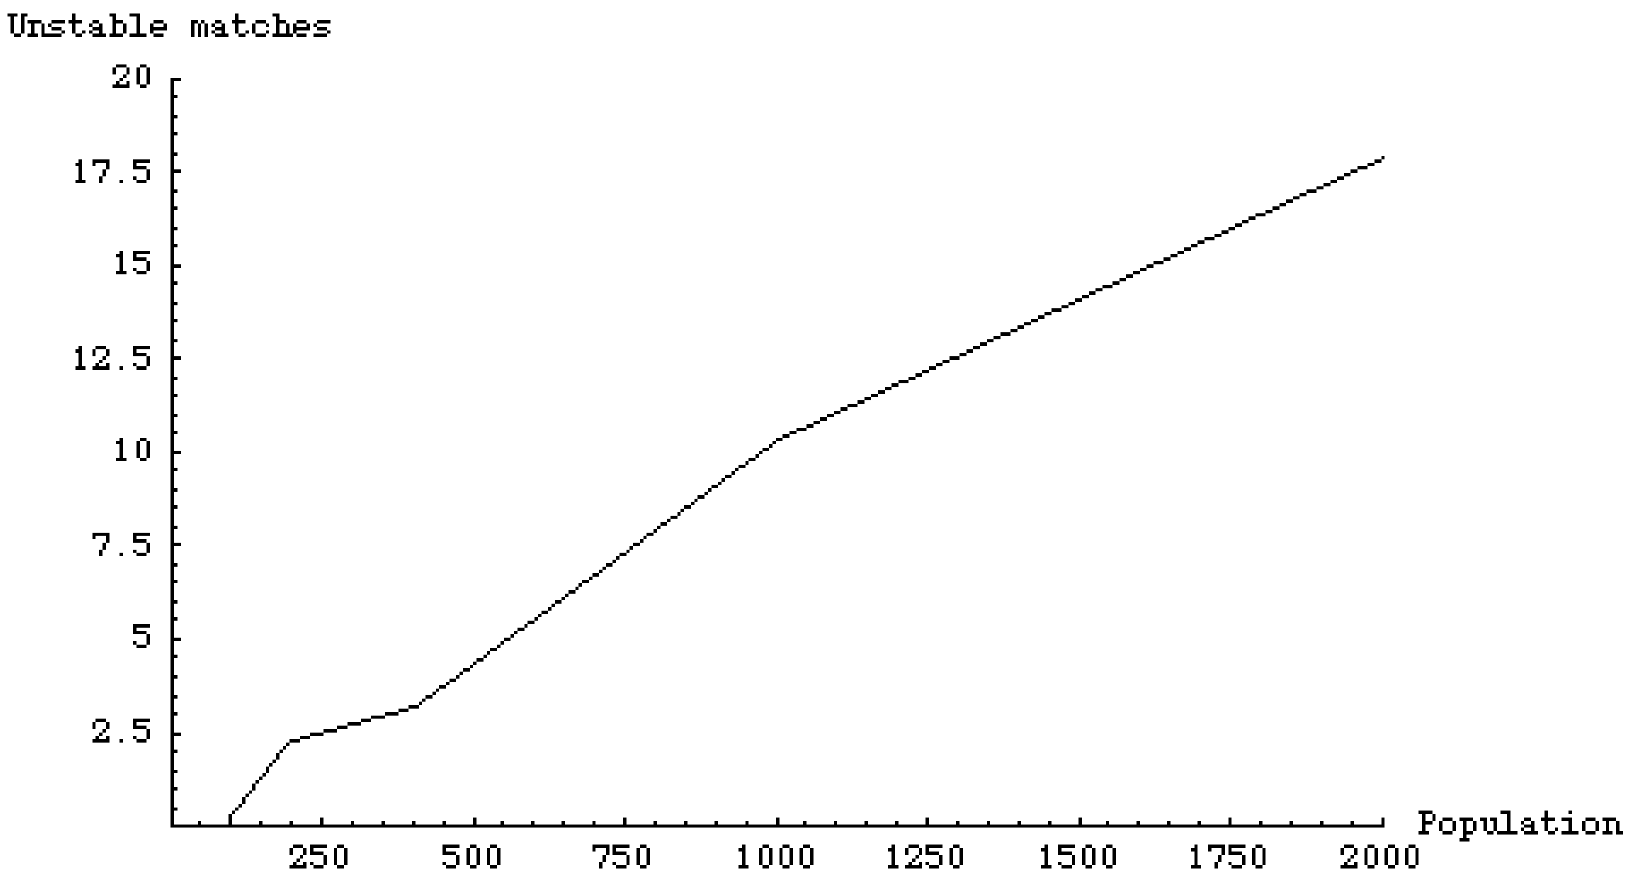

Figure 9: Figure 8 from paper. Slow increase in the number of unstable matches. Uniform random preferences, 100 periods, 1% acceptance.

Discussion

1. See the paper for more results.

2. In a nutshell: Distributed Matching is more realistic, fairer and has superior SW.

3. And can be adapted to more general settings.

4. We marvel at the strange attractiveness of equilibrium concepts.

Briefly: Simple Linear Learning

- What happens if parametric learning rules are applied in strategic contexts?
- Explored the basic model in *Stochastic Models of Learning* [Bush and Mosteller, 1955].
- Generalized the model to two-player games (learning in a strategic context).
- Implemented in `bushandmosteller1a.nlogo/html`.

The setup of [Bush and Mosteller, 1955] may be formalized as follows using stochastic linear operators for updating the action probabilities. Let:

A_i A series of r distinct actions, one of which is taken by an agent at each choice point. $i = 1, 2, \ldots, r$.

$\mathbf{p}_t$ A column vector of r elements, $p_1, p_2, \ldots, p_r$. p_i is the probability that the agent will take action A_i at the next choice point. $\mathbf{p}_0$ is the initial distribution of the action probabilities.

E_j A series of n event types, whose occurrences typically depend upon which action was taken by the agent. $j = 1, 2, \ldots, n$.

$e_{j,t}$ The event of event type E_j occurring at time t. Again: time is discrete.

$\mathbf{T}_j$ An $r \times r$ stochastic matrix (columns are probabilities adding to 1). $j = 1, 2, \ldots, n$.

Given $\mathbf{p}_t$ followed by action A_i and (crucially) event E_j, $\mathbf{p}_{t+1}$ is generated as follows:

$$\mathbf{p}_{t+1} = \mathbf{T}_j \mathbf{p}_t \tag{11}$$

$\mathbf{T}_j$ *transforms* $\mathbf{p}_t$ to yield $\mathbf{p}_{t+1}$, and does so in a linear fashion. $\mathbf{T}_j$ is said to be a *linear operator.* The full model, with n events and corresponding $\mathbf{T}_j$'s is then:

$$\mathbf{p}_{t+1} = \begin{cases} \mathbf{T}_1 \mathbf{p}_t & \text{if } E_1 \\ \mathbf{T}_2 \mathbf{p}_t & \text{if } E_2 \\ \vdots & \vdots \\ \mathbf{T}_n \mathbf{p}_t & \text{if } E_n \end{cases} \tag{12}$$

For example, if upon starting E_2 occurs, then E_3, then E_2, then E_1, we have

$$\mathbf{p}_4 = \mathbf{T}_1 \mathbf{T}_2 \mathbf{T}_3 \mathbf{T}_2 \mathbf{p}_0 \tag{13}$$

Let's look at the NetLogo implementation

`opim-sky.wharton.upenn.edu/~sok/age/nlogo/`

then

`bushandmosteller1a.html`

Part D:
Discussion & Conclusion

Discussion

- Ongoing theme in these studies: simple models of learning in strategic contexts.
- Learning as requiring risk and exploration.
- Findings include: learning as meliorization; risk may be overcome and collective benefits achieved (including tacit collusion); size matters

Conclusion: Rationality

The theory of games, indeed much of social science, has been built upon Rational Choice Theory, yet this theory has led to many unsatisfactory results. I am hardly the first to question it. Sen's critique, "Rational Fools" [Sen, 1977], remains apt. Here is recent statement of his.

> Rationality is interpreted here, broadly, as the discipline of subjecting one's choices—of actions as well as of objectives, values and priorities—to reasoned scrutiny. Rather than defining rationality in terms of some formulaic conditions that have been proposed in the literature (such as satisfying some prespecified axioms of "internal consistency of choice," or being in conformity with "intelligent pursuit of self-interest," or being some variant of maximizing behavior), rationality is seen here in much more general terms as the need to subject one's choices to the demands of reason. [Sen, 2002, page 4]

Appendix 1:

Puzzles and Problems with Classical Game Theory

From Binmore [Binmore, 1992, pages 50–1]

Here's what got me started.

> What is important here is that game theory does not pretend to tell you how to make judgments about the shortcomings [in terms of ideal rationality] of an opponent. **In making such judgments, you would be better advised to consult a psychologist than a game theorist.** Game theory is about what players will do when it is understood that both are rational in some [specific, very strong] sense. . . .

Problems, puzzles, etc. with classical game theory: Heroic assumptions

In classical game theory [Bacharach, 2006, page xiv]:

- “each player of a game [actually] maximizes [its] expected utility, given the expected behaviour of [the] other players”
- “players have common knowledge of the game itself”
- players “have common knowledge of one another’s rationality”

Can any finite machine actually have common knowledge?
Is actual maximization universally credible?

Problems, puzzles, etc. with classical game theory: Solution concepts

Nash equilibrium, standardly, but

> Except for the two-person constant-sum game, it has long been recognized that there is no unique solution concept that can be regarded as *the* natural extension of individual rational behavior to multiperson decision making. This is one of the central problems of game theory. [Shubik, 1982, page 369]

Nash equilibrium and departures from ideal rationality?

Problems, puzzles, etc. with classical game theory: Predictions

> Ironically, [classical] game theory is often hoisted on its own pétard: many of its most fundamental predictions—predictions that would have been too vague to test with any confidence in the pre-game-theoretic era—are *decisively and repeatedly disconfirmed,* in laboratory settings, with substantial agreement among experimenters, regardless of their theoretical priors. [Gintis, 2000, page xxiv]

And in experiments, even Prisoner's Dilemma, there is substantial, systematic departure from game theoretic predictions.

Problems, puzzles, etc. with classical game theory: Choosing which among many equilibria, 1

- Given multiple equilibria, which one will be chosen?
- To the extent that the theory cannot distinguish among different equilibria, it fails for lack of specificity.
- Consider the Heads and Tails (aka: Schelling) game. In experiments, $> 75\%$ of human pairs are able to coordinate [Bacharach, 2006, p. 2].

	heads		tails	
heads	1	1	0	0
tails	0	0	1	1

Problems, puzzles, etc. with classical game theory: Choosing among many equilibria, 2; the Hi-Lo game

Bacharach's Hi-Lo game[Bacharach, 2006, p. 6]:

	heads		tails	
heads	2	2	0	0
tails	0	0	1	1

> ...Hi-Lo presents a fundamental problem for game theory. From the assumptions that the players are perfectly rational (in the normal sense of maximizing expected payoffs) and that they have common knowledge of their rationality, we cannot deduce that each will choose *high*. Or, expressing the same idea in normative terms, there is no sequence of steps of valid reasoning by which perfectly rational players can arrive at the conclusion that they ought to choose *high.* Many people will find this claim incredible, but it is true. [Bacharach, 2006, p. 6]

Problems, puzzles, etc. with classical game theory: Coordinating among many equilibria

- For games with multiple equilibria distinguish:
  - the choice problem (Which equilibrium will be chosen?) and
  - the coordination problem (How can agents coordinate their plays to arrive at an equilibrium outcome?).
- In all of this the Folk Theorem exacerbates the problems: in repeated games nearly any sequence of play belongs to some Nash equilibrium.
- Coordination is a mystery in classical game theory [Kimbrough and Axtell, 2006]. See next page for an example.

	c_1	c_2
r_1	(3, 3)	(0, 2)
r_2	(2, 0)	(1, 1)

Figure 10: Standard Stag Hunt (SSH). Player R has $\Sigma^R = \{r_1, r_2\}$. Player C has $\Sigma^C = \{c_1, c_2\}$. A good label for r_1 and c_1 is 'Hunt Stag' and for r_2 and c_2, 'Hunt Hare'.

	c_1	c_2	$(c_1, \frac{1}{2}; c_2, \frac{1}{2})$
r_1	$(3, 3)$	$(0, 2)$	$(1.5, 2.5)$
r_2	$(2, 0)$	$(1, 1)$	$(1.5, 0.5)$
$(r_1, \frac{1}{2}; r_2, \frac{1}{2})$	$(2.5, 1.5)$	$(0.5, 1.5)$	$(1.5, 1.5)$

Table 3: Expected payoffs for row and column, each playing one of its strategies supported by a Nash equilibrium (could be IER). Game of Figure 10, page 73.

Appendix 2:
Two-Sided Matching
Backup Foils

Outline

"The High Cost of Stability in Two-Sided Matching: How Much Social Welfare Should be Sacrificed in the Pursuit of Stability?" by Axtell and Kimbrough `http://opim-sky.wharton.upenn.edu/~sok/sokpapers/2009/Matching-WCSS.pdf`

1. Background on two-sided matching
2. Critique of received view
3. Agent model and results
4. Discussion

1. Background on two-sided matching

- No non-arbitrary definition of the class of problems. Basic idea: two sets of items, R and C. You want to match items from R with items from C and you have some measure of value or performance for the matching. How do you do it? Well, lots of special conditions and so forth. Instead of trying to say it all abstractly, we'll begin concretely with an example, then generalize it.

- That said, there are some established definitions, which we shall use.

  Remember: "the Marriage Problem" is really a huge abstraction of the marriage problem. We can study it without thinking that it somehow captures the essence of the underlying problem.

The Marriage Problem (as traditionally described)

- There are m men and n women. Each man ranks each of the women with respect to his willingness or desire to marry her. Rank 1 is best, rank n (m) is least preferred. The women similarly rank the men.

- How should we do the matching (of each of the m men with each of the n women)?

- We need a measure of performance (MOP), some way to assess how good a matching is. Then we can look for ways to achieve a matching that does well in terms of the measure of performance.

Stability: One MOP (measure of performance)

- Suppose we have Bob, Carol, Ted and Alice. Bob and Carol are married. Ted and Alice are married. Bob prefers Alice as a wife to Carol. Alice prefers Bob as a husband to Ted. This situation is said to be unstable.

- Consider the preference orderings (rankings) shown in Figure 11.

	Carol	Alice
Bob	2, 1	1, 1
Ted	2, 2	1, 2

Figure 11: A sample array of 2-sided preference orderings

Marking marriages with $\Rightarrow \ldots \Leftarrow$

	Carol	Alice
Bob	$\Rightarrow$2, 1$\Leftarrow$	1, 1
Ted	2, 2	$\Rightarrow$1, 2$\Leftarrow$

Figure 12: Bob is married to Carol, Ted is married to Alice

Notice that under the present arrangements, Bob would prefer to be married to Alice and Alice would prefer being married to Bob. Instability!

NB: This is not a game in strategic form.

Review of the Two-Sided Matching Problem and G-S [Gale and Shapley, 1962]

In a two-sided matching problem, we are given two disjoint sets[1] R and C (think: Row and Column). Denote the elements of R as $r_1, r_2, \ldots, r_m$ and the elements of C as $c_1, c_2, \ldots, c_n$.

Definition 1. [Matched Pair] *A* **matched pair** *is an ordered 2-tuple, $\langle r, c \rangle$, for which either*

- $r \in R$ *and either (*$c \in C$ *xor* $r = c$*), or*
- $c \in C$ *and either (*$r \in R$ *xor* $r = c$*).*

[1] See [Gale and Shapley, 1962, page 12] for problems when the sets are not disjoint.

Under the intended interpretation if $r \in R$ and $c \in C$, then $\langle r, c \rangle$ is the match (e.g., marriage) of r and c. The definition, however, allows *self-matchings*, e.g., $\langle r, r \rangle$ and $\langle c, c \rangle$. A *restricted matched pair* is a matched pair formed under conditions that prohibit self matches. The interpretation of a self-match is that that particular self prefers not to be matched with any of the lower-valued possibilities. (E.g., would prefer to be unmarried.)

Definition 2. [Match Array] *A* **match array** *is an enumeration of all (legal) matched pairs.*

The following is a tabular representation in schematic form of a match array. Let $|R| = m$ and $|C| = n$.

	c_1	c_2	$\dots$	c_n	r_s
r_1	$\langle r_1, c_1\rangle$	$\langle r_1, c_2\rangle$	$\dots$	$\langle r_1, c_n\rangle$	$\langle r_1, r_1\rangle$
r_2	$\langle r_2, c_1\rangle$	$\langle r_2, c_2\rangle$	$\dots$	$\langle r_2, c_n\rangle$	$\langle r_2, r_2\rangle$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
r_m	$\langle r_m, c_1\rangle$	$\langle r_m, c_2\rangle$	$\dots$	$\langle r_m, c_n\rangle$	$\langle r_m, r_m\rangle$
c_s	$\langle c_1, c_1\rangle$	$\langle c_2, c_2\rangle$	$\dots$	$\langle c_n, c_n\rangle$	

Figure 13: Schematic representation in tabular format of a match array for R and C

The row heading c_s indicates self matches for the corresponding columns, and the column heading r_s indicates self matches for the corresponding rows. If matched pairs are unrestricted (i.e., self matching is permitted), the size of the match array is $|R| \cdot |C| + |R| + |C|$. Note that $\langle c_s, r_s \rangle$ (the southeast corner of the table) is not permitted in the match array. If restricted matching is in force the size is $|R| \cdot |C|$. The *restricted match array* is the $m \times n$ sub-table obtained from the match array by eliminating the column and row of self matches.

Definition 3. [Match] *A* **match** *(or matching)* μ *is a set of matched pairs from a match array.*

Any individual member of R or C may have a *quota*, an upper bound on the number of matches in which it may participate. In marriage problems everyone characteristically has a quota of 1. In matching workers to firms, typically each worker has a quota of 1 and firms may have quotas greater than 1.

Definition 4. [Feasible Match] *A match* μ *is a* **feasible match** *if no quota is violated.*

In terms of the match array, a row element is feasible if the number of matched pairs in the row is not greater than the row element's quota. Similarly a column element is feasible if the number of matched pairs in the column is less than or equal to its quota. A match is feasible if all of the row and column elements in its corresponding match array are feasible.

Definition 5. [Simple Match] *A* **simple match** *is a feasible match under the condition that each member of* R *or* C *has a quota of 1.*

A *valuation* for a row element $r \in R$, V_r, is a preference ranking for r on $C \cup r$. Similarly for a valuation V_c of a column element c. We assume that ranking is strict (no ties) and stipulate that the most preferred object has a rank of 1 and the least preferred object the rank of $|C| + 1$ ($|R\| + 1$). (In a restricted match, with no self-matching, the maximum rank values are $|R|$ and $|C|$.) Denote the rank (preference index) of y for x by $P_x y$ (mnemonic: P is *position* in the ranking, lower numbered being better). For each matched pair $\langle x, y\rangle$ there is a corresponding *valuation pair*, $\langle P_x y, P_y x\rangle$. A *valuation array* (or "ranking matrix" [Gale and Shapley, 1962, page 11]) is a match array with valuation pairs substituted for matched pairs.

	c_1	c_2	$\ldots$	c_n	r_s
r_1	$\langle P_{r_1}c_1, P_{c_1}r_1\rangle$	$\langle P_{r_1}c_2, P_{c_2}r_1\rangle$	$\ldots$	$\langle P_{r_1}c_n, P_{c_n}r_1\rangle$	$\langle P_{r_1}r_1, P_{r_1}r_1\rangle$
r_2	$\langle P_{r_2}c_1, P_{c_1}r_2\rangle$	$\langle P_{r_2}c_2, P_{c_2}r_2\rangle$	$\ldots$	$\langle P_{r_2}c_n, P_{c_n}r_2\rangle$	$\langle P_{r_2}r_2, P_{r_2}r_2\rangle$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
r_m	$\langle P_{r_m}c_1, P_{c_1}r_m\rangle$	$\langle P_{r_m}c_2, P_{c_2}r_m\rangle$	$\ldots$	$\langle P_{r_m}c_n, P_{c_n}r_m\rangle$	$\langle P_{r_m}r_m, P_{r_m}r_m\rangle$
c_s	$\langle P_{c_1}c_1, P_{c_1}c_1\rangle$	$\langle P_{c_2}c_2, P_{c_2}c_2\rangle$	$\ldots$	$\langle P_{c_n}c_n, P_{c_n}c_n\rangle$	

Table 4: Schematic representation in tabular format of a valuation array for R and C

Definition 6. [Acceptability] *A match is* **unacceptable** *if it contains a matched pair for which $P_x y > P_x x$. A match is* **acceptable** *if it is not unacceptable.*

If x is matched to y in μ (write: $\mu(x) = y$) and x prefers to be self-matched to being matched to y, then μ is unacceptable. Put otherwise, in an acceptable match, no one does worse than being self-matched.

Definition 7. [Individually Rational Matching] *A match μ is* **individually rational** *if it is feasible and acceptable.*

Comment: this definition corresponds to and is equivalent to Definition 2.2 of Roth and Sotomayor [Roth and Sotomayor, 1990, page 21].

We are now in position to be able to define the class of two-sided match problems, given R and C and the definitions above.

NB: Broader, alternative definitions are possible. This has thus a "technical" meaning.

Definition 8. [Two-Sided Match Problem] *A* **two-sided match problem** *is the problem of finding a high-quality, ideally optimal, individually rational matching between the two given sets R and C, and a given valuation array.*

Two important special cases of two-sided match problems are the marriage problem and the admissions problem.

Definition 9. [Marriage Problem] *A* **marriage problem** *is a two-sided match problem in which each member of $R \cup C$ has a quota of 1.*

Definition 10. [Admissions Problem] *An* **admissions problem** *is a two-sided match problem in which at least one member of R (C, think: colleges) has a quota greater than 1 and all members of C (R, think: students) have a quota of 1.*

We will for the most part focus our attention on marriage problems.

Note that either problem may allow or disallow self-matches. By default we assume that self-matches are allowed and that in any matching if an agent is not matched to an agent of the opposite class (an r is matched to some member of C, and a c is matched to some member of R), then the agent is matched to itself.

Obtaining any significant results for matching problems requires that "high-quality" in the definition of two-sided matching be specified. Of course, the notions of Pareto optimal matches (matches such that no other match exists that makes everyone better off) and Pareto dominant matches (matches such that in every other match everyone is strictly worse off) are available to us. The literature on two-sided matching has, however, focused on the concept of a stable match, which may be defined as follows.

Definition 11. [Stable Match] *A match* μ *is* **unstable** *if it contains matched pairs* $\langle w, x \rangle$, $\langle y, z \rangle$ *such that* $P_w x > P_w z$ *and* $P_z y > P_z w$. *A match is* **stable** *if it is not unstable.*

Note: Valuations with lower rank values are preferred to valuations with higher rank values, so if $P_w x > P_w z$, then w prefers z over x. This definition corresponds to Gale and Shapley's definition [Gale and Shapley, 1962, page 10] and to Definition 2.3 of Roth and Sotomayor [Roth and Sotomayor, 1990, page 21].

Stability characterizes equilibrium in two-sided matching problems. Given a match there may well be many individuals with incentive to find different partners. Characterizing two-sided matching problems is the assumption that individuals cannot act alone. If r_i is unhappy with its partner in a match, we assume that r_i can do nothing unless there is a c_j that r_i prefers to its current partnet, and c_j is less happy with its match than it would be paired with r_i. The marriage metaphor motivates the definition. If a man prefers the wife of his neighbor to his own wife, and the neighbor's wife prefers the man to her own husband, we have an unstable situation. A matching is thus said to be stable if there is no pair of matched couples for which a pair of individuals has incentive to switch partners. Note that if a switch is made the jilted individuals may or may not benefit.

The literature on two-sided matching has largely been directed at finding stable matches. The seminal paper is [Gale and Shapley, 1962], which presents an algorithms that are guaranteed to find stable matches for marriage problems and admissions problems. Originally referred to as the *deferred acceptance procedure*, the algorithms collectively has come to be called the Gale-Shapley (G-S) procedure. When necessary, we shall distinguish them by name: G-S$_M$ for the marriage version (everyone has a quota of 1), and G-S$_A$ for the admissions version. We will get to the G-S procedure shortly.

Gale and Shapley present two theorems in their paper.

Theorem 1. [G-S Theorem 1 [Gale and Shapley, 1962]] *There always exists a stable set of marriages.*

That is, for the marriage problem, the G-S$_M$ algorithm (deferred acceptance) will always find a stable match.

Assuming a definition of optimality—"A stable assignment is called *optimal* if every applicant [in an admissions problem; applicants have quotas of 1] is at least as well off under it as under any other stable assignment"—Theorem 2 asserts that G-S$_A$ produces an optimal match.

Theorem 2. [G-S Theorem 2 [Gale and Shapley, 1962]] *Every applicant is at least as well off under the assignment given by the deferred acceptance procedure as he would be under any other stable assignment.*

Description of G-S$_M$

“To start, let each boy propose to his favorite girl. Each girl who receives more than one proposal rejects all but her favorite from among those who have proposed to her. However, she does not accept him yet, but keeps him on a strong to allow for the possibility that someone better may come along later.

“We are now ready for the second stage. Those boys who were rejected now propose to their second choices. Each girl receiving proposals chooses her favorite from the group consisting of the new proposers and the boy on her string, if any. She rejects all the rest and again keeps the favorite in suspense.

“We proceed in the same manner. Those who are rejected at the second stage propose to their next choices, and the girls again reject all but the best proposal they have had so far.

"Eventually (in fact, in at most $n^2 - 2n + 2$ stages) every girl will have received a proposal, for as long as any girl has not been proposed to there will be rejections and new proposals, but since no boy can propose to the same girl more than once, every girl is sure to get a proposal in due time. As soon as the last girl gets her proposal the 'courtship' is declared over, and each girl is now required to accept the boy on her string." [Gale and Shapley, 1962, pages 12–3]

2. Critique of received view

- So what's not to like?
- There are anomalies or paradoxical results (here we proceed as philosophers)

Anomaly #1: Sexual Bias (fairness)

G-S in terms of our match array representation, figure 13, is row-oriented. It matters whether men are assigned to the set R or whether women are. To illustrate, on page 11 they present their *Example* 1, in the form of a valuation array for a 3×3 marriage problem.

Example 1. The following is the "ranking matrix" [valuation array] of three men, α, β, and γ, and three women, A, B, and C.

	A	B	C
α	1,3	2,2	3,1
β	3,1	1,3	2,2
γ	2,2	3,1	1,3

The first number of each pair in the matrix gives the ranking of women by the men, the second number is the ranking of the men by the women. Thus, α ranks A first, B second, C third, while A ranks β first, γ second, and α third, etc.

There are six possible sets of marriages; of these, three are stable. One of these is realized by giving each man his first choice, thus α marries A, β marries B, and γ marries C. Note that although each woman gets her last choice, the arrangement is nevertheless stable. Alternatively one may let the women have their first choices and marry α to C, β to A, and γ to B. The third stable arrangement is to give everyone his or her second choice and have α marry B, β marry C, and γ marry A. The reader will easily verify that all other arrangements are unstable.

Transparently, in this example, for whichever stable match of the three stable matches is found by the G-S it will be the case that there is another stable match in which someone does worse than in the match found by G-S. If G-S finds the first match, then all the men do worse in matches 2 and 3. If G-S finds the second match, then all the men fare better in the first and third matches. If G-S finds the third match, then all the men do better in the first match. Optimality is defined by G-S as Pareto dominance. There is here no Pareto dominant solution among the stable matches.

If the men are assigned to R in this example, then G-S finds the match in which each man gets his first choice. Contrariwise, if the women get the the R rôle, the G-S finds the match in which each woman gets here first choice. G-S cannot find the stable match in which everyone gets his or her second choice. This, as shown in [Gale and Shapley, 1962], is not a peculiarity of the example. Under the G-S$_M$ procedure, whichever group gets assigned to be R obtains a matching that is optimal for it in the following sense:

Definition 12. [From [Gale and Shapley, 1962]] *A stable assignment is called* **optimal** *if every applicant is at least as well off under it as under any other stable assignment.*

See also Theorem 2.12 of [Roth and Sotomayor, 1990, page 32]. And in general (assuming throughout strict preferences), the side assigned R gets its optimal stable matching, while the side assigned C gets its pessimal stable matching; each member of C gets its least preferred achievable partner ([Roth and Sotomayor, 1990, page 33], [Knuth, 1976]).

Anomaly #2: G-S Theorem 2 Revisited (fairness)

The sexual bias results for G-S$_M$ raise the question of social welfare. Under G-S one class is assigned to be R and in consequence gets results that are, in the sense defined, optimal for it. The other class gets its worst deal. Although there may exist compromise matches, they are not available via G-S, which is also silent on which class to favor.

Theorem 2 of [Gale and Shapley, 1962], the proof of optimality for the G-S$_A$ algorithm (deferred acceptance) for the admissions problem, might seem to avoid the social welfare question. It is important to understand that it does not. The algorithm simply *defines* the applicants to be the favored class, at the expense of the universities. To illustrate, consider the following problem. There are two schools, Hard University and Knocks College, and each has a quota of 2 applicants. There are four applicants, Bob, Ted, Carol, and Alice. The valuation array (G-S's ranking matrix) is as in figure 14.

	Hard	Knocks
Bob	2,1	1,4
Ted	1,2	2,3
Carol	2,3	1,2
Alice	1,4	2,1

Figure 14: Valuation array for an admissions problem

As always we have two cases for the G-S algorithm, depending on which side is assigned R (and hence does the proposing) and which C. Here is the G-S description of deferred acceptance for the admissions problem when the students propose (are Row players) [Gale and Shapley, 1962, page 13]:

First, all students apply to the college of their first choice. A college with a quota of q then places on its waiting list the q applicants who rank highest, or all applicants if there are fewer than q, and rejects the rest. Rejected applicants then apply to their second choice and again each college selects the top q from among the new applicants and those on its waiting list, puts these on its new waiting list, and rejects the rest. The procedure terminates when every applicant is either on a waiting list or has been rejected by every college to which he is willing and permitted to apply. At this point each college admits everyone on its waiting list and the stable assignment has been achieved.

Case 1: Students propose. In the example to hand, in the first round Bob and Carol apply to Knocks, and Ted and Alice apply to Hard. At this point the procedure terminates. Figure 15 shows the pairings; it is obvious the match is stable.

	Hard	Knocks
Bob		1,4
Ted	1,2	
Carol		1,2
Alice	1,4	

Figure 15: Stable match with students proposing

What happens in case 2, with the colleges proposing? Gale and Shapley [Gale and Shapley, 1962] simply do not propose an algorithm for admissions with the colleges proposing. G-S for admissions problems stipulates that students, not colleges, propose. Can there be stable matches other than the one found by G-S? Figure 16 is one for the example to hand.

	Hard	Knocks
Bob	2,1	
Ted	1,2	
Carol		1,2
Alice		2,1

Figure 16: Second stable matching for the example

Is the matching of figure 16 obtainable by a close analog of G-S$_A$, but with colleges proposing (assigned to R)? Consider the following procedure.

Procedure 1. [Colleges Proposing Admissions] *All colleges propose to their highest-ranked students. A student with more than one proposal rejects all but its highest-ranked proposal. Any college with a presently unrejected proposal adds the student to its waiting list, decrementing* *its quota. Every college with a positive quota remaining proposes to its next most preferred student. Every student with more than one proposal rejects all but its highest-ranked proposal. Every college now updates it quota to reflect the number of its accepted proposals. The procedure terminates when every college has proposed to each of its ranked students, or before if stopping will not change the outcome.*

We can now describe case 2, using the Colleges Proposing Admissions procedure.

Case 2: Schools propose. Hard proposes to Bob and Knocks proposes to Alice. There are no conflicts so round 2 ensues. Hard proposes to Ted and Knocks proposes to Carol. Again, there are no conflicts. The procedure terminates because both schools have proposed to all of their ranked students. The resulting match is that of figure 16.

As you may easily verify, the Colleges Proposing Admissions is directly analogous to the G-S$_A$ procedure, and generally shares its properties, with the exception that the colleges get their optimal match.

Anomaly #3: Social Welfare and Unstable Matches

Definition 13. [Summation Social Welfare Function] *In a two-sided matching problem, with sides* X *and* Y*, the summation social welfare of a matching,* μ*, is the sum total of rank points realized by each side in* μ*.*

Definition 14. [SumProduct Social Welfare Function] *In a two-sided matching problem, with sides* X *and* Y*, the sumproduct social welfare of a matching,* μ*, is the sum total of the products of rank points of the pairs that are matched in* μ*.*

Proposition 1. [Existence of Socially Superior Unstable Matches] *In the two-sided matching problem with* $|X| > 3$ *and* $|Y| > 3$*, there exist individual preference orderings such that there are stable matches that are socially inferior to unstable matches, under the summation and sumproduct social welfare functions.*

Summing Up

There are other anomalies/problems as well. Key for us is to note that there are (at least) three criteria by which to evaluate a matching:

1. Stability (G-S et al., [Roth and Sotomayor, 1990])

2. Social welfare (NB linear programming formulation)

3. Fairness (e.g., [Klaus and Klijn, 2006], [Fuku et al., 2006])

Stable matchings under G-S do not maximize social welfare and do not lead to fair outcomes.

3. Agent model and results

The challenge:

How might we find matchings that either (a) a stable and are improved wrt social welfare and/or fairness, or (b) are not stable, but not unacceptably so, and are improved wrt social welfare and/or fairness?

How, in particular, might we build agent-based models to investigate the problems of two-sided matching?

NB: Real-world complications often vex the deferred-acceptance approach.

Read the paper, but briefly. . .

Agent Model: Distributed Matching

- Two types of agents.
- Each individual of each type randomly (uniformly) ranks each individual of the other type.
- Each step a small number of agents, randomly chosen, become active and make proposals of engagement that are accepted or not.
- Steps accumulate into periods, when the number of agents activated equals the number of agents.
- With some probability, an engaged pair get married.

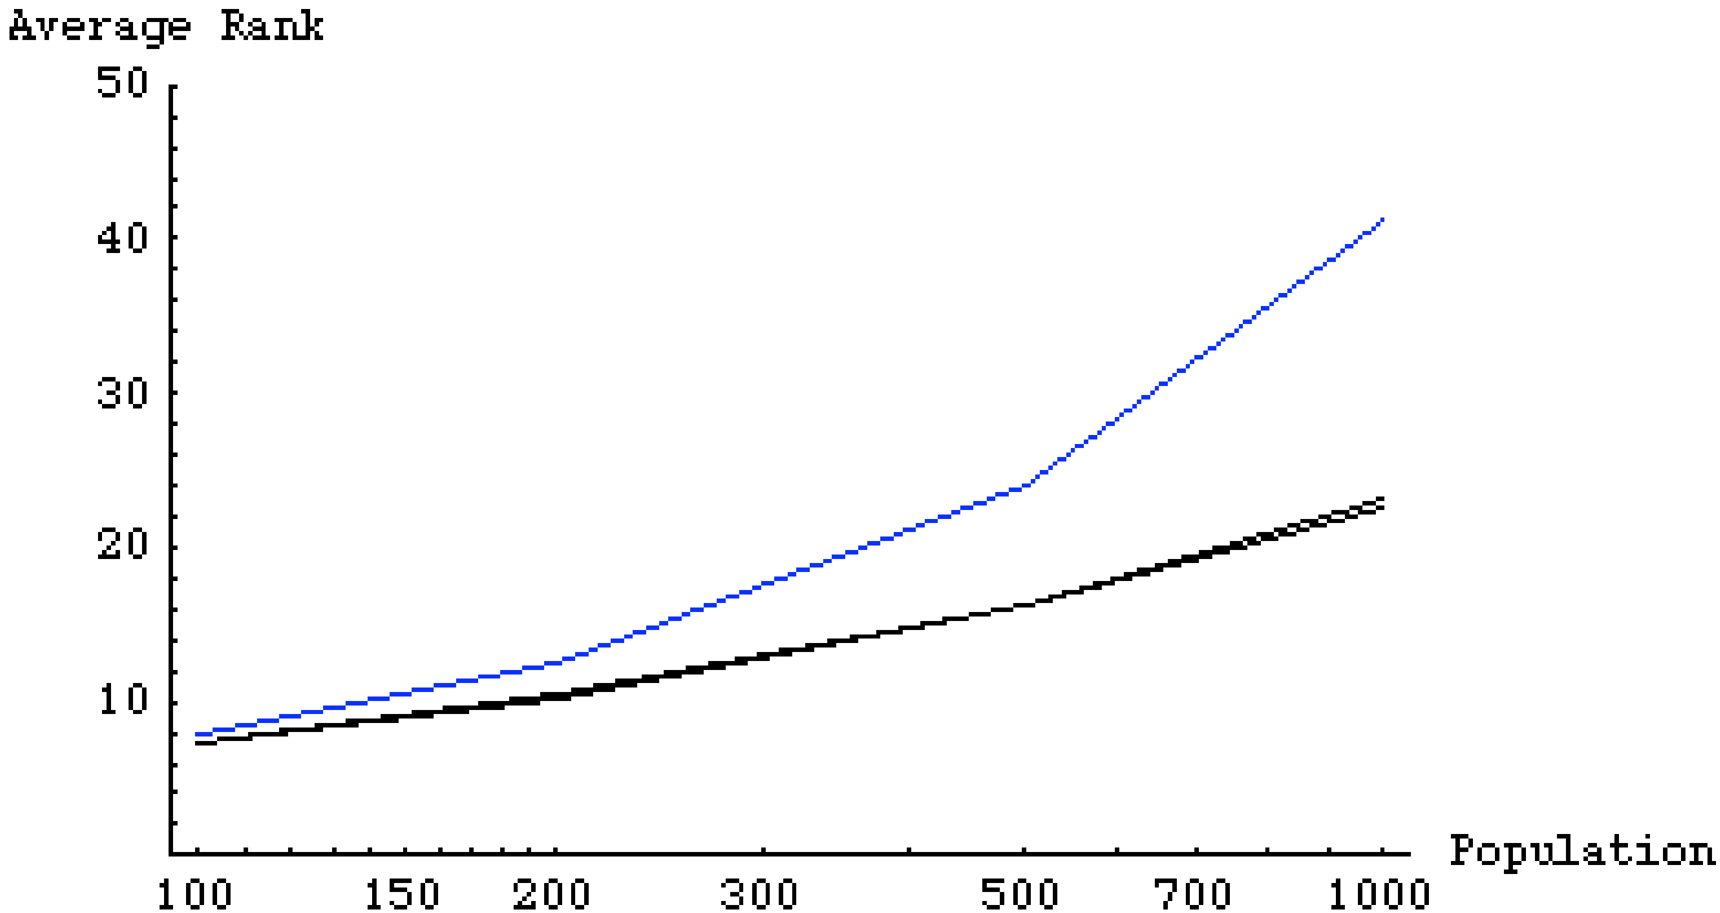

Figure 17: Figure 4 from paper. Compared to G-S (blue), our agents are better on SW and on fairness.

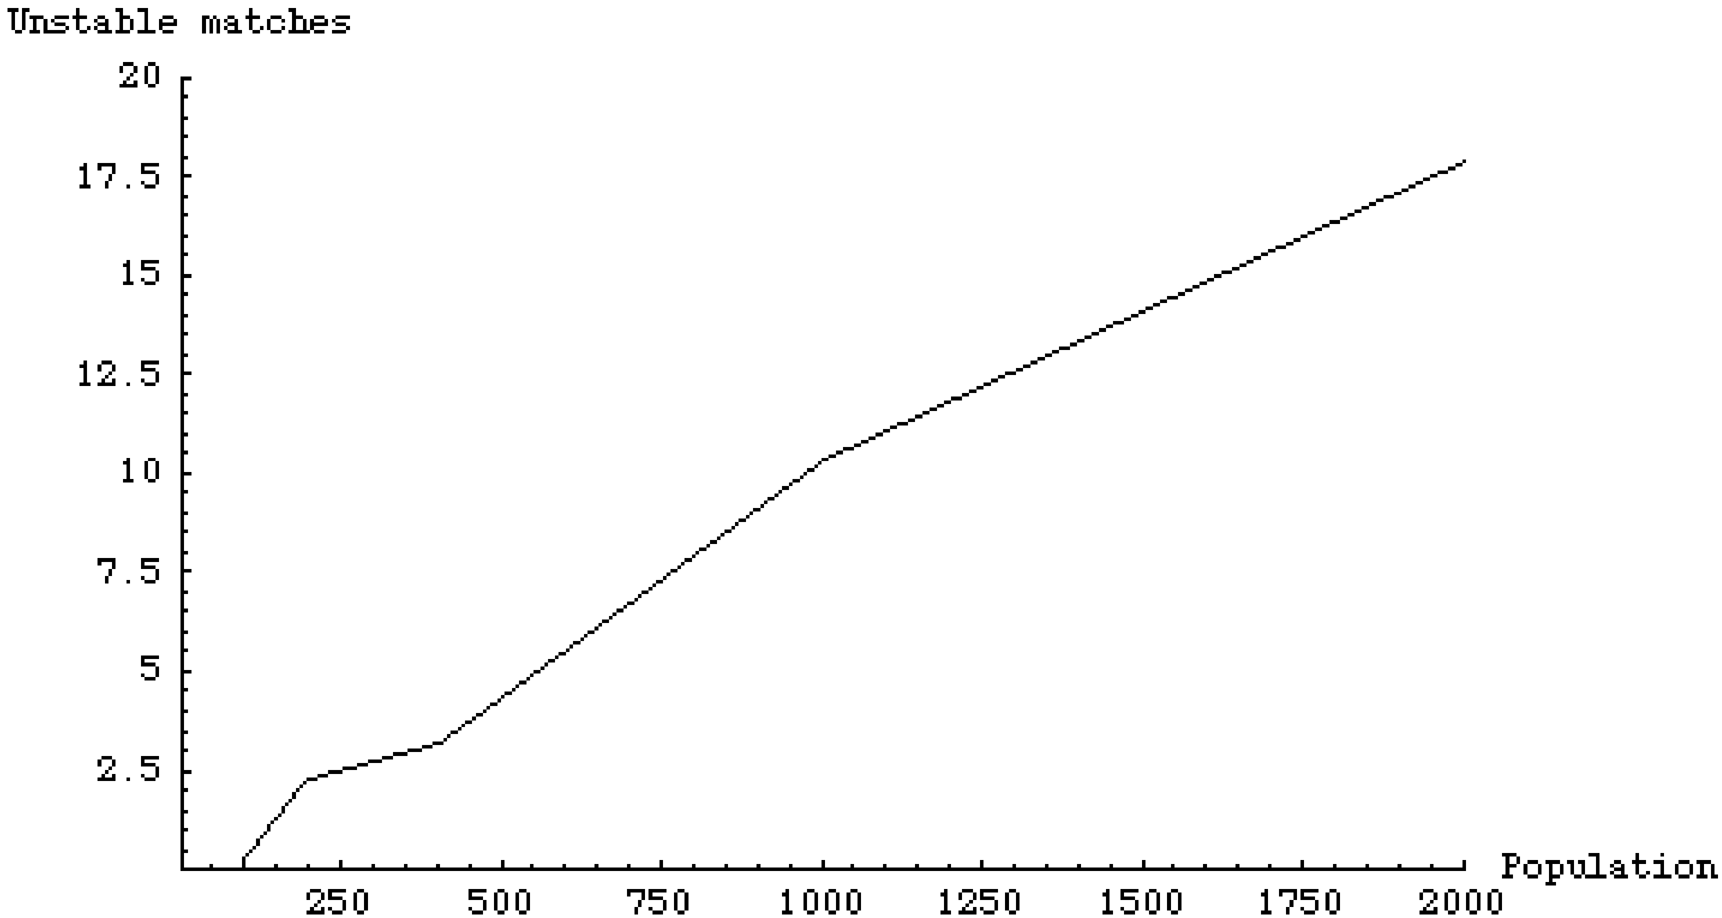

Figure 18: Figure 8 from paper. Slow increase in the number of unstable matches. Uniform random preferences, 100 periods, 1% acceptance.

4. Discussion

1. See the paper for more results.
2. In a nutshell: Distributed Matching is more realistic, fairer and has superior SW.
3. And can be adapted to more general settings.
4. We marvel at the strange attractiveness of equilibrium concepts.

References

[Axtell and Kimbrough, 2008] Axtell, R. L. and Kimbrough, S. O. (2008). The high cost of stability in two-sided matching: How much social welfare should be sacrificed in the pursuit of stability? In *Proceedings of the 2008 World Congress on Social Simulation (WCSS-08)*. `http://mann.clermont.cemagref.fr/wcss/`.

[Bacharach, 2006] Bacharach, M. (2006). *Beyond Individual Choice*. Princeton University Press, Princeton, NJ. edited by Natalie Gold and Robert Sugden.

[Binmore, 1992] Binmore, K. (1992). *Fun and Games: A Text on Game Theory*. D.H. Heath and Company, Lexington, MA.

[Bush and Mosteller, 1955] Bush, R. R. and Mosteller, F. (1955). *Stochastic Models for Learning*. Wiley, New York, NY.

[Fuku et al., 2006] Fuku, T., Namatame, A., and Kaizouji, T. (2006). Collective efficiency in two-sided matching. In Mathieu, P., Beaufils, B., and Brandouy, O., editors, *Artificial Economics: Agent-Based Methods in Finance, Game Theory and Their Applications*, pages 115–126. Springer-Verlag, Berlin, Germany.

[Gale and Shapley, 1962] Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

[Gintis, 2000] Gintis, H. (2000). *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press, Princeton, NJ.

[Kimbrough, 2008] Kimbrough, S. O. (2008). On artificial intelligence, evolutionary computation and networks. In *Proceedings of the Wharton-INSEAD Alliance Conference on Network-based Strategies and Competencies*.

[Kimbrough and Axtell, 2006] Kimbrough, S. O. and Axtell, R. L. (2006). On concepts of rationality in games. In *Society for Exact Philosophy*. http://opim-sky.wharton.upenn.edu/~sok/sokpapers/2006/two-concepts-rationality.pdf.

[Kimbrough and Lu, 2005] Kimbrough, S. O. and Lu, M. (2005). Simple reinforcement learning agents: Pareto beats Nash in an algorithmic game theory study. *Information Systems and e-Business Management*, 3(1):1–19. http://dx.doi.org/10.1007/s10257-003-0024-0.

[Kimbrough and Murphy, 2009] Kimbrough, S. O. and Murphy, F. H. (2009). Learning to collude tacitly on production levels by oligopolistic agents. *Computational Economics*, 33(1):47–78. `http://dx.doi.org/10.1007/s10614-008-9150-6`.

[Klaus and Klijn, 2006] Klaus, B. and Klijn, F. (2006). Procedurally fair and stable matching. *Economic Theory*, 27:431–447.

[Knuth, 1976] Knuth, D. E. (1976). *Marriages Stables*. Les Presses de l'Université de Montreal, Montreal, Canada.

[Roth and Sotomayor, 1990] Roth, A. E. and Sotomayor, M. A. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, Cambridge, United Kingdom.

[Sen, 2002] Sen, A. (2002). *Rationality and Freedon*, chapter

Introduction: Rationality and Freedom, pages 3–64. Harvard University Press, Cambridge, MA.

[Sen, 1977] Sen, A. K. (1977). Rational fools: A critique of the behavioural foundations of economic theory. *Philosophy and Public Affairs*, 6:317–344.

[Shubik, 1982] Shubik, M. (1982). *Game Theory in the Social Sciences*. The MIT Press, Cambridge, MA.

[Varian, 2003] Varian, H. R. (2003). *Intermediate Microeconomics: A Modern Approach*. W. W. Norton & Company, New York, NY, sixth edition.

End Note

$Id: comprats-master-foils.tex 633 2009-01-23 12:55:25Z sok $