




EDITED BY
NIR VULKAN, ALVIN E. ROTH
& ZVIKA NEEMAN

the handbook of
MARKET
DESIGN

OXFORD

THE HANDBOOK OF
MARKET DESIGN

This page intentionally left blank

THE HANDBOOK OF

MARKET DESIGN

Edited by

NIR VULKAN,
ALVIN E. ROTH,

and

ZVIKA NEEMAN

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Oxford University Press 2013

The moral rights of the authors have been asserted

First Edition published in 2013

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data
Data available

Library of Congress Control Number: 2013944501

ISBN 978-0-19-957051-5

Printed and bound by
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work.

We would like to dedicate this book to Fiona, Norette and Tom Vulkan; Emilie, Aaron and Ben Roth; and Haya and Uri Neeman.

This page intentionally left blank

CONTENTS

.....

<i>List of Figures</i>	x
<i>List of Tables</i>	xiii
<i>List of Contributors</i>	xvi

Introduction	1
NIR VULKAN, ALVIN E. ROTH, AND ZVIKA NEEMAN	

PART I GENERAL PRINCIPLES

1. What Have We Learned From Market Design?	7
ALVIN E. ROTH	
2. Not Up To Standard: Stress Testing Market Designs for Misbehavior	51
GARY E. BOLTON	
3. Using and Abusing Auction Theory	62
PAUL KLEMPERER	

PART II CASES

SECTION II.A MATCHING MARKETS

4. Market Design for Kidney Exchange	93
TAYFUN SÖNMEZ AND M. UTKU ÜNVER	
5. School Choice	138
ATILA ABDULKADİROĞLU	
6. Improving Efficiency in School Choice	170
AYTEK ERDİL AND HALUK ERGİN	
7. Can the Job Market for Economists Be Improved?	189
SARBARTHA BANDYOPADHYAY, FEDOR ISKHAKOV, TERENCE JOHNSON, SOOHYUNG LEE, DAVID MCARTHUR, JOHN RUST, JOEL WATSON, AND JOHN WATSON	
8. Designing Markets for Ideas	222
JOSHUA S. GANS AND SCOTT STERN	

9. Redesigning Microcredit 249
ASHOK RAI AND TOMAS SJÖSTRÖM

SECTION II.B AUCTIONS

10. The Product-Mix Auction: A New Auction Design for Differentiated Goods 269
PAUL KLEMPERER
11. Optimal Incentives in Core-Selecting Auctions 282
ROBERT DAY AND PAUL MILGROM
12. Auctioning Rough Diamonds: A Competitive Sales Process for BHP Billiton's Ekati Diamonds 299
PETER CRAMTON, SAMUEL DINKIN, AND ROBERT WILSON

SECTION II.C E-COMMERCE

13. Ending Rules in Internet Auctions: Design and Behavior 325
AXEL OCKENFELS AND ALVIN E. ROTH
14. Designing Markets for Mixed Use of Humans and Automated Agents 345
ANDREW BYDE AND NIR VULKAN
15. The Design of Online Advertising Markets 363
BENJAMIN EDELMAN
16. Very-Large-Scale Generalized Combinatorial Multi-Attribute Auctions: Lessons from Conducting \$60 Billion of Sourcing 379
TUOMAS SANDHOLM
17. Designing Automated Markets for Communication Bandwidth 413
NIR VULKAN AND CHRIS PREIST

SECTION II.D LAW DESIGN

18. A Mechanism Design Approach to Legal Problems 437
ALON KLEMENT AND ZVIKA NEEMAN
19. Legislation with Endogenous Preferences 456
AVIAD HEIFETZ, ELLA SEGEV, AND ERIC TALLEY

PART III EXPERIMENTS

- | | | |
|-----|--|-----|
| 20. | Common-Value Auctions with Liquidity Needs: An Experimental Test of a Troubled-Assets Reverse Auction | 489 |
| | LAWRENCE M. AUSUBEL, PETER CRAMTON, EMEL FILIZ-OZBAY, NATHANIEL HIGGINS, ERKUT Y. OZBAY, AND ANDREW STOCKING | |
| 21. | Information Disclosure in Auctions: An Experiment | 555 |
| | MARTIN DUFWENBERG AND URI GNEEZY | |
| 22. | Experiments with Buyer-Determined Procurement Auctions | 573 |
| | ELENA KATOK | |
| 23. | The Inefficiency of Splitting the Bill | 587 |
| | URI GNEEZY, ERNAN HARUVY, AND HADAS YAFE | |

PART IV COMPETING DESIGNS

- | | | |
|-----|--|-----|
| 24. | Competing Mechanisms | 609 |
| | MICHAEL PETERS | |
| 25. | Three Case Studies of Competing Designs in Financial Markets | 629 |
| | NIR VULKAN AND ZVIKA NEEMAN | |
| | <i>Index</i> | 661 |

LIST OF FIGURES

1.1	Potential kidney exchanges between patient–donor pairs at multiple centers	19
1.2	High-level flowchart of the Roth and Peranson (1999) applicant-proposing deferred acceptance algorithm with couples	38
4.1	A three-way kidney exchange	95
4.2	Graphical representation for the set of recipients in example 1	104
4.3	AB–O type pair saving two underdemanded pairs in a three-way exchange	109
4.4	Overdemanded pairs B–O / AB–A each saving one underdemanded pair and an A–B type pair in a three-way exchange	109
4.5	An overdemanded AB–O type pair can save three underdemanded pairs in a four-way kidney exchange	110
4.6	A three-way list exchange chain	130
4.7	A simultaneous three-way altruistic donor chain	130
4.8	A non-simultaneous two-way altruistic donor chain	131
4.9	All feasible exchanges between three pairs at two centers	134
7.1	Growth in EJM ads, applicants, recruiters and recommenders	199
10.1	An example of bids in the Bank of England’s auction	273
10.2	Equilibrium in the Bank of England’s auction	275
12.1	Price index for rough diamonds, 2007–13	301
12.2	Forming aggregate demand from individual demands	306
12.3	Price versus quantity in a uniform-price auction	306
12.4	Two bidders, blue and red, in a uniform-price auction	307
12.5	Price versus quantity in a pay-as-bid auction	308
12.6	Typical bid curves in pay-as-bid and uniform-price auctions	308
12.7	Price versus quantity in an ascending-clock auction	310
13.1	Cumulative distributions over time of auctions’ last bids	326
13.2	Late bidding to protect information	331
13.3	Late bidding as best response to incremental bidding	332
13.4	Percentage of bidders who snipe (bids in stage 2), and linear trends	337

13.5	“Early” (final stage-1) prices as percentage of final price and linear trends	338
15.1	Deceptive banner ads overstate the urgency of clicking through	365
16.1	A relatively simple discount schedule	382
16.2	A simple example of bidding with alternates, cost drivers, attributes, and constraints	383
16.3	An example of bidding with cost structures and attributes	384
16.4	A user interface for expressive allocation evaluation by the bid taker	388
16.5	Advanced sourcing application platform	399
16.6	A tradeoff query in our system	405
16.7	A comparison query in our system in a medical sourcing event	405
18.1	The mechanism design problem expressed as a Mount–Reiter triangle	443
20.1	Difference between value and price in a four-bidder sealed-bid auction	517
20.2	Bidder profit on security sales in a four-bidder sealed-bid auction	518
20.3	Subject Payoffs (sales plus bonus) in a four-bidder clock auction	519
20B.1	Probability density and cumulative distribution of common value	529
20C.1	Probability density of common value in the more precise case, by size of holdings	539
20C.2	Probability density of common value in the less precise case, by size of holdings	539
20D.1	Probability density and cumulative distribution of common value	544
20E.2	Probability density of common value by size of holdings	552
21.1	Average bids and winning bids, (a) session F1, (b) session F2, (c) session S1, (d) session S2, (e) session N1, (f) session N2	566
21.2	Average winning bids	568
21.3	Cumulative distributions of the bids chosen in period 10 for each treatment, aggregated across the two sessions for each treatment	569
22.1	Average number of bids bidders place in the two open-bid treatments	579
22.2	Average prices in the Katok and Wambach (2008) study	580
22.3	A risk-neutral incumbent’s bidding functions when $K = 2$ and $K = 20$	583
22.4	Bidding behavior: bids as a function of x_i , proportion of $Bid = R$, proportion of $Bid \geq \max(x_i, K)$	584
23.1	Summary of the restaurant results	595
23.2	Cumulative distribution graphs for treatments 1–3 of the restaurant experiment	596
25.1	Determinants of trade mechanism selection in derivatives markets	635

25.2	BT rules are stricter (and thus BT contracting costs higher) at North American derivatives exchanges than in European or Japan	636
25.3	Options variables on different exchanges, 2002–05	639
25.4	Futures variables on different exchanges, 2002–05	640
25.5	Contracts variables on different exchanges, 2002–05	641
25.6	Frequency of large trades as a function of market liquidity	643
25.7	Standardized average trade size for SETS and dealer market trade, 1998–2002	646
25.8	Percentage value through SETS vs. percentage value at SETS price	648
25.9	Combinations of trade sizes that the SETS and the SEAQ mechanisms can select	650
25.10	Volatility of the various sectors studied by Henry (2003)	651
25.11	Percentage of market share of orders, NYSE versus ECNs	654
25.12	Price-correlated market share in orders of 100–49 shares, by industry (NYSE versus ECNs)	654
25.13	Price-correlated market share in orders of 5000–9999 shares, by industry (NYSE versus ECNs)	656

LIST OF TABLES

4.1	Patient and living-donor distributions used in simulations	113
4.2	Simulation results for the average number of patients actually matched and predicted by the formulas to be matched	116
4.3	A Pareto-efficient two-way exchange mechanism outcome for n pairs	132
5.1	Properties of the mechanisms	148
5.2	Welfare consequences of tie breaking and strategy proofness for grade 8 applicants in NYC in 2006–07	156
7.1	Location and employment status of EJM candidates	201
7.2	Distribution of candidates' primary fields	202
7.3	Characteristics of EJM job advertisements	203
7.4	Distributions of advertised research fields	204
7.5	Applications submitted by applicants	205
7.6	Average number of applications submitted by primary field of applicant	206
7.7	Applications received by recruiters	207
7.8	Recommenders, recommendees and recommendations	207
12.1	BHP Billiton term auction, February 2009	320
14.1	Summary of the experiments	359
14.2	Summary of buyers' performance as measured by the average price of transactions	360
14.3	Summary of buyers' performance as measured by the percentage of the target quantity purchased	360
15.1	Payment rules for different formats of online advertising	365
19.1	Benchmark case vs. maximizing wealth case	474
19.2	Benchmark case vs. maximizing happiness case	474
20.1	Schedule of treatments	496
20.2	Holdings of securities by bidder and security	501
20.3	Comparison of mean outcomes by auction type in Experiment 1 with liquidity bonus	504

20.4	Comparison of mean outcomes by auction type in Experiment 1 without liquidity bonus	505
20.5	Comparison of mean outcomes by auction type in Experiment 2 with liquidity bonus	505
20.6	Comparison of mean outcomes by auction type in Experiment 2 without liquidity bonus	506
20.7	Experiment 1: experimental subject fixed effects	507
20.8	Experiment 1: Effect of competition and expected payoff	509
20.9	Experiment 1: Effect of session-specific expected payoff	510
20.10	Experiment 1: Fixed effects regressions (fully interacted) with liquidity needs	511
20.11	Experiment 1: Fixed effects regressions (fully interacted) without liquidity needs	512
20.12	Experiment 2: pooled-security reverse auction	513
20.13	Experiment 2: Payoffs over time	513
20.14	Experiment 2: Effect of the lowest applied bidder strategy ratio on outcomes	515
20C.1	Holdings of securities by bidder and security	535
20E.1	Holdings of securities by bidder and security	549
21.1	The bids in session F1 (full information feedback)	559
21.2	The bids in session F2 (full information feedback)	560
21.3	The bids in session S1 (semi-information information feedback)	561
21.4	The bids in session S2 (semi-information feedback)	562
21.5	The bids in session N1 (no information feedback)	563
21.6	The bids in session N2 (no information feedback)	564
21.7	A pairwise comparison of bids in the first period across sessions	568
22.1	Actual and predicted buyer surplus levels in the Engelbrecht-Wiggans et al. (2007) study	575
22.2	Average buyer surplus levels, proportion of efficient allocations, and the comparison between actual and estimated theoretical buyer surplus levels and efficiency	578
22.3	Summary of the average prices, bid decrements and theoretical predictions in the Elmaghraby et al. (2012) study	582
23.1	Summary of the restaurant results	594
23.2	Hypothesis tests on the restaurant results: p -values for hypotheses 1–3	595
23.3	Summary of the results relevant to hypothesis 4	597

23.4	The p -values for hypothesis 4	598
23.5	The t -test p -values (two-tailed) for the null hypothesis of no gender effects in the restaurant results	599
23.6	Production table for a laboratory examination of the diner's dilemma	600
23.7	Results of the laboratory experiment	601
23B.1	The individual revenue and cost from each quantity in the laboratory experiment	604
25.1	Overall daily rank scores for volatility on both the dealer market and SETS	647

LIST OF CONTRIBUTORS

Atila Abdulkadiroğlu is Professor in Economics at Duke University. He taught at Northwestern University and Columbia University before coming to Duke. He received his PhD in Economics at the University of Rochester. He has consulted school districts in redesigning student assignment systems, including Boston (MA), Chicago (IL), Denver (CO), New Orleans (LA), and New York City (NY). He is a recipient of an Alfred P. Sloan Research Fellowship and a National Science Foundation CAREER award. He serves as an Editor-in-Chief of *Review of Economic Design* and on the board of the Institute for Innovation in Public School Choice.

Lawrence M. Ausubel is Professor of Economics at the University of Maryland. He has published widely on auctions, bargaining, the credit card market, and other aspects of industrial organization and financial markets. He has been awarded fifteen US patents relating to auction methodology and he has several other patents pending. He received his AB in Mathematics from Princeton University, his MS in Mathematics from Stanford University, his MLS in Legal Studies from Stanford Law School, and his PhD in Economics from Stanford University. He is also chairman of Power Auctions LLC and president of Market Design Inc.

Sarbartha Bandyopadhyay manages the servers and database systems for EconJob-Market.org. He is also the maintainer of the EJM codebase. He is the Co-President of Technoluddites Inc. and a developer for Editorial Express, Conference Maker and HeadHunter.

Gary E. Bolton is the O.P. Jindal Chair of Management Measurement Economics in the Jindal School of Management, University of Texas at Dallas, and is Director of the Laboratory of Behavioral Operations and Economics. He studies economic and business decision-making and strategic games, with special interest in bargaining, cooperation, reputation building, social utility, and strategic learning.

Andrew Byde is Head of Research at Acunu Ltd, a database startup, where his research focuses on algorithms and data structures for data storage systems. Prior to joining Acunu, he spent ten years at Hewlett-Packard laboratories, publishing in diverse areas of computer science, including autonomous agents, biologically inspired computing and market-based design. He held a Royal Society Industry Fellowship to the University of Southampton for two years. His PhD is in Mathematics, from Stanford University.

Peter Cramton is Professor of Economics at the University of Maryland. Since 1983, he has conducted widely cited research on auction theory and practice. The main focus is the design of auctions for many related items. Applications include auctions for radio spectrum, electricity, financial securities, diamonds, and timber. He has introduced innovative market designs in many industries. He has advised numerous governments on market design and has advised dozens of bidders in major auction markets. He received his BS in Engineering from Cornell University and his PhD in Business from Stanford University.

Robert Day is currently an Associate Professor of Operations and Information Management at the University of Connecticut. He received his PhD in Applied Mathematics with a concentration in Operations Research from the University of Maryland in 2004. His dissertation, which focused on combinatorial auctions, received INFORMS' Dantzig dissertation award in 2005. He continues to study combinatorial auctions and other related applications, including markets for grid computing, and the scheduling of operating-room resources in order to improve hospital efficiency. Further, he has recently consulted on the design of combinatorial auctions for spectrum licenses held in the UK and other countries.

Samuel Dinkin is Senior Auction Consultant at Power Auctions LLC, where he helps governments and companies design, implement, and participate in high-stakes auctions on six continents. Since 1995, he has designed and conducted over fifty auctions totaling over \$100 billion in transactions, including auction rules for the world's four largest multi-unit electricity auctions and first multi-unit telecom auctions for seven countries. He is a seventh-plateau IBM master inventor. He captained the US contract bridge team, winning a silver medal in the 2009 world championships. He received his BS in economics from Caltech and his MA and PhD from the University of Arizona.

Martin Dufwenberg got his PhD in Uppsala in 1995; he worked at universities in Tilburg, Uppsala, Stockholm, Bonn, and is currently affiliated with the University of Arizona, the University of Gothenburg and Bocconi University. His research uses game theory and experiments to incorporate insights from psychology into economic analysis.

Benjamin Edelman is an Associate Professor at the Harvard Business School. His research explores the public and private forces shaping Internet architecture and business opportunities, including online advertising, consumer protection, and regulation. He is a member of the Massachusetts Bar. His writings are available at <<http://www.benedelman.org>>.

Aytek Erdil is a Lecturer in Economics at the University of Cambridge, and is a fellow of King's College, Cambridge. He received his PhD from the University of Chicago, and held postdoctoral fellowships at Harvard Business School and University of Oxford, prior to his current position.

Haluk Ergin is Associate Professor in the Department of Economics at the University of California, Berkeley.

Emel Filiz-Ozbay is Assistant Professor of Economics at the University of Maryland. She received her PhD in Economics from Columbia University in 2007 and joined the University of Maryland faculty in the same year. She is interested in experimental and behavioral economics, decision theory, industrial organization, and market design. She has conducted several auction and market design experiments to better understand the performance of different mechanisms. Her research also addresses the behavioral biases of agents in their decision-making, and how those biases vary between agents and depending on the environment. She teaches courses in microeconomics and contract theory.

Joshua S. Gans is a Professor of Strategic Management and holder of the Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship at the Rotman School of Management, University of Toronto. While his research interests are varied, he has developed specialties in the nature of technological competition and innovation, economic growth, publishing economics, industrial organization and regulatory economics. In 2007, he was awarded the Economic Society of Australia's Young Economist Award, and in 2008 was elected as a Fellow of the Academy of Social Sciences, Australia.

Uri Gneezy is The Epstein/Atkinson Chair in Behavioral Economics and Professor of Economics & Strategy at Rady School of Management, University of California San Diego. As a researcher, his focus is on putting behavioral economics to work in the real world, where theory can meet application. He is looking for basic research as well as more applied approaches to such topics as incentives-based interventions to increase good habits and decrease bad ones, "pay-what-you-want" pricing, and the detrimental effects of small and large incentives. In addition to the traditional laboratory and field studies, he is currently working with several firms, conducting experiments in which basic findings from behavioral economics are used to help companies achieve their traditional goals in non-traditional ways.

Ernan Haruvy is an Associate Professor in Marketing at the University of Texas at Dallas. He earned his PhD in Economics in 1999 from the University of Texas at Austin. His main interests are in market design applications, including auctions, procurement, matching, learning, e-commerce, and software markets. His main methodological tools come from behavioral and experimental economics. He publishes in various disciplines with publications in journals such as *American Economic Review*, *Marketing Science*, *Journal of Marketing Research* and *Journal of Finance*.

Aviad Heifetz is a Professor of Economics at the Economics and Management Department of the Open University of Israel, where he served as department chair, 2006–09. He is the author of the textbook *Game Theory: Interactive Strategies in Economics and Management* (Cambridge University Press, 2012). His research in game theory and economic theory provided insights into the evolution of preferences, market design,

bargaining, competitive economies with asymmetric information, and interactive epistemology. He serves on the editorial boards of *Games and Economic Behavior*, *International Journal of Game Theory*, and *Mathematical Social Sciences*.

Nathaniel Higgins is an Economist at the Economic Research Service, United States Department of Agriculture (USDA), and Adjunct Professor in the Department of International Economics at the Johns Hopkins School of Advanced International Studies. He uses experimental and computational economics to study problems in market design, especially problems related to the design of USDA conservation programs. He has published articles on the design of auctions, behavioral economics, and commodities prices, and consulted in the design of auctions for spectrum and airport landing slots. He received his BA in Mathematics and Economics from Ithaca College and his PhD in Agricultural and Resource Economics from the University of Maryland.

Fedor Iskhakov is Senior Research Fellow at ARC Centre of Excellence in Population Ageing Research at University of New South Wales. He received his PhD in Economics from the University of Oslo, Norway, in 2009 and Candidate of Sciences degree from St Petersburg State University, Russia, in 2006. Iskhakov's PhD dissertation, "A dynamic structural analysis of health and retirement," was awarded the HM King of Norway golden medal (H.M. Kongens gullmedalje) as best research in social sciences among young researchers in Norway in 2008.

Terence Johnson is Assistant Professor of Economics at the University of Notre Dame. He received his PhD from the University of Maryland in 2011, specializing in microeconomic theory and industrial organization.

Elena Katok is Ashbel Smith Professor at the Naveen Jundal School of Management, at the University of Texas at Dallas. She has co-authored a number of scholarly articles in behavioral operations management, focusing on using laboratory experiments to test game-theoretic models of contracting and competitive procurement. She has been a member of INFORMS since 1995, and was the winner of the 2000 Franz Edelman competition. She is one of the organizers of the annual Behavioral Operations Management conference, a Department Editor for Behavioral Operations, and a Deputy Editor at the *Production and Operation Management (POM) Journal*, and the President of the INFORMS Section for Behavioral Operations Management.

Alon Klement writes on various subjects in the fields of Civil Procedure, Law and Economics and Law and Social Norms. He teaches at the Radzyner School of Law, in the Interdisciplinary Center, Herzliya. In recent years he has visited and taught at Columbia University and Boston University in the US, and at the University of Bologna, in Italy. He earned his LLB in Law and BA in Economics from Tel Aviv University. After practicing law for several years in a private law firm in Israel he went to Harvard, where he earned his SJD degree.

Paul Klemperer is the Edgeworth Professor of Economics at Oxford University. He has advised numerous governments, including devising the UK government's 3G

mobile-phone license auction, which raised £22.5 billion, and assisting the US Treasury in the financial crisis. He has also developed new auction designs; his most recent innovation—the product-mix auction—is regularly used by the Bank of England. He co-invented the concept of “strategic complements;” developed the “supply function” analysis of electricity markets, and the theory of consumer switching costs; and has applied techniques from auction theory in a range of other economic contexts, from finance to political economy.

Soohyung Lee is an Assistant Professor at University of Maryland. She received her PhD from Stanford University and BA from Seoul National University. Prior to starting her PhD program, she served the Ministry of Strategy and Finance in Korea, as a Deputy Director from 1999 to 2002. Her research interests broadly lie in applied econometrics and market design.

David McArthur is an Economist at the Federal Reserve Board of Governors. His research interests are information economics, networks, and industrial organization. His PhD is in Economics from the University of Maryland.

Paul Milgrom is the Leonard and Shirley Ely Professor of Humanities and Sciences at Stanford University and a member of the National Academy of Sciences and American Academy of Arts and Sciences. He has published widely on the subject of auctions and market design, including thirty published articles, a book (*Putting Auction Theory to Work*) and several issued and pending patents. He is inventor or co-inventor of some of the leading auction methods used for large auctions today, including the simultaneous multiple-round auction, the clock-proxy auction, the core-selecting auction, and the assignment auction. He is the founder and chairman of Auctionomics and leads its team of economists assisting the FCC in designing and implementing its “incentive auction” to buy television broadcast rights and repurpose the spectrum for wireless broadband services.

Zvika Neeman is Professor of Economics at the Berglas School of Economics at Tel Aviv University. He is a microeconomic and game theorist with diverse interests who specializes in mechanism design. Before joining Tel Aviv University, he held positions at Boston University and at the Hebrew University of Jerusalem. He received his PhD from Northwestern University in 1995.

Axel Ockenfels is Professor of Economics at the University of Cologne, Director of the Cologne Laboratory of Economic Research, and Coordinator of the DFG research group “Design & Behavior.” He publishes in leading journals in economics, but also in business administration, information systems, psychology, and sociology, as well as in application-oriented outlets. Ockenfels is a Member of the Berlin-Brandenburgische and of the North Rhine-Westphalian Academy of Sciences, the Academic Advisory Board at the Federal Ministry of Economics and Technology, and the Scientific Advisory Board of the University of Cologne. In 2005 he received the Gottfried Wilhelm Leibniz Prize of the German Science Foundation.

Erkut Y. Ozbay is Assistant Professor of Economics at the University of Maryland. He received his PhD in Economics from New York University in 2007 and joined the faculty at the University of Maryland in the same year. He is also the Director of the Experimental Economics Laboratory at the University of Maryland. His research interests are experimental economics and theory. His research mainly focuses on understanding how economic agents make decisions when they are faced with different types of uncertainty and how their behavior is affected by their experience, concerns, and the change of their understanding of the economic problem they are facing.

Michael Peters is Professor of Economics at the University of British Columbia since 2003 and a Fellow of the Econometric Society. His research focuses on search and competing mechanisms.

Chris Preist is Reader in Sustainability and Computer Systems at University of Bristol Computer Science Department. He is also a member of the Cabot Institute and the Systems Centre. His research interests include life cycle analysis of digital services with particular reference to the digital transformation of the news and media industry. Prior to joining Bristol, he was a master scientist at HP Labs, Bristol. He was the technical lead on the EU Framework 5 Semantic Web-based Web Services project, coordinating input from over twenty researchers across eight institutions, and chair of the Scientific Advisory Board of the Framework 6 DIP project. His work in the application of artificial intelligence techniques to automated diagnosis led to the deployment of several systems within HP manufacturing facilities and the development of an award-winning product (Agilent Fault Detective). He has a degree in Mathematics from Warwick University, and a PhD in the Semantics of Logic Programming from Imperial College, London.

Ashok Rai is a development economist who works on microfinance. His field research has taken him to Bangladesh, Colombia, India, and Kenya. Rai has a PhD from the University of Chicago and an undergraduate degree from Stanford. He is currently an Associate Professor at Williams College in the United States, and a Member of the Courant Center at the University of Göttingen in Germany.

Alvin E. Roth received his BS from Columbia University in 1971 and PhD from Stanford University in 1974. He taught at the University of Illinois, 1974–82, at the University of Pittsburgh, 1982–98, at Harvard University, 1998–2012, and now teaches at Stanford University. He shared the 2012 Nobel Memorial Prize in Economics for his work in market design.

John Rust is Professor of Economics at Georgetown University. He was previously a Professor of Economics at University of Maryland, Yale University, and University of Wisconsin. He received his PhD from MIT in 1983, specializing in applied econometrics and computational economics. He is a co-founder of EconJobMarket.org, and Techno-luddites, Inc., which provides web-based software to assist academics in publishing, holding conferences, and evaluation of job candidates. He received the Ragnar Frisch

Medal from the Econometric Society in 1992 for his 1987 *Econometrica* paper “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.”

Tuomas Sandholm is Professor in the Computer Science Department at Carnegie Mellon University. He has published over 450 papers on market design and other topics in computer science, operations research, and game theory. He holds sixteen patents on market design. He is best known for his work on combinatorial auctions. Applications include sourcing, TV and Internet display advertising, sponsored search, and radio spectrum. He is Founder, President, and CEO of Optimized Markets, Inc. Previously he was Founder, Chairman, and CTO/Chief Scientist of CombineNet, Inc. His technology also runs the US-wide kidney exchange. He serves as the design consultant of Baidu’s sponsored search auctions. He has also consulted for Yahoo!, Netcyclor, Google, and many other companies and government institutions on market design. He is recipient of the NSF Career Award, the inaugural ACM Autonomous Agents Research Award, the Alfred P. Sloan Foundation Fellowship, the Carnegie Science Center Award for Excellence, and the Computers and Thought Award. He is Fellow of the ACM and AACL.

Ella Segev is a researcher in the Department of Industrial Engineering and Management at Ben-Gurion University, Beer Sheva, Israel. Her research interests include auction theory, bargaining theory, and contests. She has a PhD from Tel Aviv University and was a research scholar at the Institute for Advanced Study in Princeton, USA. She has published papers in journals such as *International Economic Review*, *Games and Economic Behavior*, and *Public Choice*, among others.

Tomas Sjöström did his undergraduate studies in Stockholm and received a PhD from the University of Rochester. He taught at Harvard and Penn State before moving to Rutgers in 2004, where he is currently Professor of Economics. His interests include mechanism design, theories of conflict, and neuroeconomics.

Tayfun Sönmez is a Professor at Boston College, Department of Economics. He received his PhD in Economics in 1995 from the University of Rochester.

Scott Stern is School of Management Distinguished Professor and Chair of the Technological Innovation, Entrepreneurship and Strategic Management Group at the Sloan School at MIT. He explores how innovation—the production and distribution of ideas—differs from more traditional economic goods, and the implications of these differences for entrepreneurship, business strategy, and public policy. He received his PhD from Stanford University, and he is the Director of the Innovation Policy Working Group at the National Bureau of Economic Research. In 2005 he was awarded the Kauffman Prize Medal for Distinguished Research in Entrepreneurship.

Andrew Stocking is a market design economist at the US Congressional Budget Office (CBO). He provides analysis of market rules and their effect on expected outcomes for markets, which include environmental cap-and-trade programs, spectrum auctions, Medicare auctions, oil and gas markets, and financial markets. Prior to working at the

CBO, he designed and worked with online advertising auctions, charitable fundraising markets, airport auctions for landing slots, and international telecom auctions. He has published several papers on the unintended consequences of market rules. He holds a BS in chemical engineering and an MS in environmental engineering, both from Stanford University, and a PhD in resource economics from the University of Maryland.

Eric Talley is the Rosalinde and Arthur Gilbert Professor in Law, Business and the Economy at the University of California Berkeley (Boalt Hall) School of Law. His research focuses on the intersection corporate law, firm governance and financial economics. He holds a bachelor's degree from UC San Diego, a PhD in economics from Stanford, and a JD also from Stanford, where he was articles editor for the *Stanford Law Review*. He has previously held permanent or visiting appointments at the University of Southern California, the University of Chicago, Harvard University, Georgetown University, the California Institute of Technology, Stanford University, the University of Sydney (Australia), and the University of Miami. He is a frequent commentator on the radio show *Marketplace*, and often speaks to corporate boards and regulators on issues pertaining to fiduciary duties, governance, and corporate finance.

M. Utku Ünver is a Professor at Boston College, Department of Economics. He received his PhD in Economics in 2000 from the University of Pittsburgh.

Nir Vulkan is Economics Professor at the Said Business School and a Fellow of Worcester College, both at Oxford University. He is the authors of dozens of articles on market design and the book *The Economics of e-Commerce* (Princeton University Press, 2003). He has worked with many software and e-commerce companies designing markets mainly on the Internet, which are used by humans and software agents. His algorithms for automated trading have been used by hedge funds to trade futures in markets all over the world.

Joel Watson is Professor of Economics at the University of California, San Diego. His research mainly addresses how contractual relationships are formed and managed, and the role of institutions, using game-theoretic models. He authored a popular textbook on game theory (*Strategy: An Introduction to Game Theory*). He co-founded and serves as the CEO of Econ Job Market Inc., a non-profit charitable (501c3) corporation that manages application materials in the economics PhD job market. He obtained his BA from UCSD and his PhD from Stanford's Graduate School of Business. He was a Prize Research Fellow at Oxford's Nuffield College.

John Watson is Director of Analytics for the Institute of Evidence-Based Change, Encinitas, and technologist for Watson Education. He designs data collection and analysis applications primarily in the education field, oversees the largest K-20 student records database in California, and conducts research in the area of intelligent data systems. He holds a patent for mobile data-environment technologies, and has contributed to a half-dozen technology-related inventions. He received a joint doctoral degree from Claremont Graduate University and San Diego State University. He also

received degrees from University of California at San Diego (BA) and San Diego State University (MA).

Robert Wilson is an Emeritus Professor at the Stanford Business School. His main research interest is game theory, but he has worked on market designs since the early 1970s, including government auctions of exploration leases and spectrum licenses, and wholesale markets for natural gas and for electric power, reserves, and capacity. His book *Nonlinear Pricing*, in print since 1993, won the Melamed Prize of the Chicago Business School. He is a Fellow of the Econometric Society and the American Economic Association, a member of the National Academy of Sciences, and on the board of Market Design Inc.

Hadas Yafe is CEO of GreenHands.

INTRODUCTION

NIR VULKAN, ALVIN E. ROTH,
AND ZVIKA NEEMAN

“MARKET design” is the term used to refer to a growing body of work that might also be called *microeconomic engineering* and to the theoretical and empirical research that supports this effort and is motivated by it.

Economists often look at markets as given, trying to make predictions about who will do what and what will happen in these markets. Market design, in contrast, does not take markets as given; instead, it combines insights from economic and game theory together with common sense and lessons learned from empirical work and experimental analysis to aid in the design and implementation of actual markets. In recent years the field has grown dramatically—partly because of the successful wave of spectrum auctions in the US and in Europe, partly because of the clearinghouses and other marketplaces which have been designed by a number of prominent economists, and partly because of the increased use of the Internet as the platform over which markets are designed and run. There are now a large number of applications and a growing theoretical literature, which this book surveys.

Market design is both a science and an art. It is a science in that it applies the formal tools of game theory and mechanism design and it is an art because practical design often calls for decisions that are beyond the reliable scientific knowledge of the field, and because the participants in these markets are often different than they are modeled by these theories. Nevertheless, as the book demonstrates, lessons can be learned from successful and unsuccessful market designs which can be transferred to new and different environments.

In this book we attempt to bring together the latest research and provide a relatively comprehensive description of applied market design as it has taken place around the world over the last two decades or so. In particular we survey many matching markets: These are environments where there is a need to match large two-sided populations of agents such as medical residents and hospitals, law clerks and judges, or patients and kidney donors, to one another. Experience shows that if the arranged match is not appropriately stable, then participants will try to transact outside of the indicated

marketplace, and the market will unravel leading to very inefficient results. We also survey a number of applications related to electronic markets and e-commerce: The Internet is now the preferred platform for many markets and this raises some interesting issues, such as the impact of automation (for example you use a software agent to bid in an Internet auction). Also related is the resulting competition between exchanges—since anyone can access the Internet anywhere in the world, the geographic location of a market is less relevant and participants now often face a real choice of trading mechanisms which they can use. While many of the chapters in the book consider a single marketplace that has established such a dominant share of the market that most participants have no other desirable choice (e.g. medical residents), a number of chapters in this book consider the implications to market designers of the fact that participants have a choice.

Market design involves the specification of detailed rules, which are typically analyzed using what used to be called “noncooperative” game theory. The analysis focuses on the incentives for individual behavior in the particular environment considered and its consequences. Specific environments and problems can be very different from one another, and, as we’ll see, details and differences can be of huge importance in practical design. But there are also some general themes beginning to emerge from all this detail and diversity, and it will help to keep some of these in mind.

Specifically, a marketplace or the setting in which market design is performed, is part of a broader economic environment in which potential participants also have other choices to make, which may be less well known and harder to model. That is, a marketplace being designed or studied is typically part of a larger game that cannot be modeled in detail with the same confidence as the marketplace. So, to work well and attract wide participation, it may be desirable for marketplaces to promote outcomes that are in the *core* of the larger game, so that there don’t exist any coalitions that might prefer to transact outside of the marketplace, instead of participating in it.¹

A related, less formal organizing theme is that, if a marketplace is to be successful, the rules and behavior in the marketplace, together with the (unmodeled) opportunities and behavior outside the marketplace, have to form an equilibrium in which, given how the marketplace works, it makes sense for participants to enter it and participate. In this respect, experience suggests we can start to diagnose whether a marketplace is working well or badly, by examining how well it provides *thickness*, deals with *congestion*, and makes it *safe* and *simple* to participate (cf. Roth, Chapter 1).

¹ The core and various related notions of stability not only capture a very general notion of what constitutes a competitive outcome, they also apply to the less detailed models of what used to be called “cooperative” game theory, and in doing so tell us something about the options that may be available to coalitions of players even when we don’t know their strategies in detail. This is why the former distinction between cooperative and noncooperative game theory is not very useful in market design; both perspectives are employed together, to answer different kinds of question and to deal with different kinds of design constraint.

A market provides thickness when it makes many potential transactions available at the same time, so that relevant offers can be compared. (Availability in this sense has a big information component; offers must be available in a way that allows comparison.)

A market is congested if there is insufficient time or resources to fully evaluate all the potentially available transactions. Sometimes this will involve the physical resources needed to carry out transactions (e.g. they may be time consuming, and other possibilities may disappear while a transaction is being attempted), but it can also involve the information needed to make the comparisons among alternative transactions that are needed to choose among them. Congestion is thus a particular problem of thick markets with many quite heterogeneous matching opportunities, and one task of an effective market is to deal with congestion in a way that allows the potential benefits of thickness to be achieved.²

To be thick, a marketplace must also make it safe to participate, at least relative to transacting outside the marketplace. Depending on the information and sophistication of the participants, safety may also involve what kinds of strategies the rules of the marketplace require participants to be able to execute, and how sensitive it is to how well others execute their strategies. This is one of the ways in which market design differs most clearly from the theoretical literature on mechanism design, in which different mechanisms are compared by comparing their equilibria. In practical markets, particularly new ones in which all participants will begin without experience, the risks to participants out of equilibrium must also be considered, and so designers often analyze “worst cases” as well as equilibria. Unlike the presumptions made in the literature on theoretical mechanism design and implementation, market designers never know the whole game and therefore need to be cognizant of the fact that their design is one piece of a larger game. Market designers typically do not try to design a market *all* of whose equilibria accomplish something, but rather try to design a marketplace with a good equilibrium, and then try to achieve that equilibrium. If unanticipated behavior develops, the market can be modified, for example with appeals processes, or with making bidders use dropdown menus instead of typing in their own bids, and so on.

This brings us to simplicity, which involves both the market rules themselves, and the kind of behavior they elicit. Simplicity of rules is sometimes discussed under the heading of “transparency,” which also involves participants being able to audit the outcome and verify that the rules were followed. But rules may be simple and transparent yet require complex strategizing by the participants. Strategic complexity is often the more important issue, since it may affect both participation in the market, for example if implementing good strategies is costly, and market performance, by leading to mistakes and misjudgments. And the risk associated with such mistakes and misjudgments may also deter participation.

² Congestion sometimes manifests itself as coordination failure, and so signaling and other attempts to facilitate sorting are one way to deal with it. Another reaction to congestion is unraveling, i.e. starting to transact before the opening of the marketplace, and therefore often not participating in the thick market.

This volume includes chapters that provide a conceptualization of new markets or marketplaces and other designs, together with chapters that describe the adoption and implementation of specific designs (and their subsequent adjustments in light of experience), as well as the theoretical and empirical questions raised in the process. We begin with three chapters that discuss general principles in market design: Al Roth's chapter reviews some of the markets that he, his students, and colleagues have designed, and draws general conclusions from these; Gary Bolton's chapter describes how to stress test models in the lab; and Paul Klemperer's explains how to sensibly use economic theory to create good designs, and he demonstrates how using too much theory can be bad.³

Part II is the main part of the book and it provides many cases and applications of market design, some that have been running for years, and some that are still in very early stages. Part II is subdivided into sections on matching markets, auctions, e-commerce applications, and law design (a small section).

Part III focuses on market design experiments, and finally Part IV discusses the implications for market design when there is competition between markets.

³ Klemperer's chapter focuses on the design of large-scale auctions. However, we believe his advice is very relevant to all kinds of market design.

P A R T I

GENERAL
PRINCIPLES

This page intentionally left blank

CHAPTER 1

WHAT HAVE WE LEARNED FROM MARKET DESIGN?

ALVIN E. ROTH¹

INTRODUCTION

In the centennial issue of the *Economic Journal*, I wrote (about game theory) that

the real test of our success will be not merely how well we understand the general principles that govern economic interactions, but how well we can bring this knowledge to bear on practical questions of microeconomic engineering. (Roth, 1991a)

Since then, economists have gained significant experience in practical market design. One thing we learn from this experience is that transactions and institutions matter at a level of detail that economists have not often had to deal with, and, in this respect, all markets are different. But there are also general lessons. The present chapter considers some ways in which markets succeed and fail, by looking at some common patterns we see of market failures, and how they have been fixed.

This is a big subject, and I will only scratch the surface, by concentrating on markets my colleagues and I helped design in the last few years. My focus will be different than in Roth (2002), where I discussed some lessons learned in the 1990s. The relevant parts of

¹ The first part of this chapter was prepared to accompany the Hahn Lecture I delivered at the Royal Economic Society meetings, on April 11, 2007, and was published as Roth (2008a). The present chapter extends the 2008 paper with a Postscript to bring it up to date, and to include some details appropriate to this *Handbook*. I have also updated references and added some footnotes to the first part of the chapter, but otherwise it remains essentially as published in 2008. One reason for keeping this format, with a distinct Postscript to bring it up to date is that it will become clear that some of the developments anticipated in the 2008 paper have been realized in the intervening years. The work I report here is a joint effort of many colleagues and coauthors. I pay particular attention here to work with Atila Abdulkadiroğlu, Muriel Niederle, Parag Pathak, Tayfun Sönmez, and Utku Ünver. I've also benefited from many conversations on this topic with Paul Milgrom (including two years teaching together a course on market design). In the Postscript I also report on work done with Itai Ashlagi. This work has been supported by grants from the NSF to the NBER.

that discussion, which I will review briefly in the next section, gathered evidence from a variety of labor market clearinghouses to determine properties of successful clearinghouses, motivated by the redesign of the clearinghouse for new American doctors (Roth and Peranson, 1999). Other big market design lessons from the 1990s concern the design of auctions for the sale of radio spectrum and electricity; see for example Cramton (1997), Milgrom (2000), Wilson (2002), and, particularly, Milgrom (2004).²

As we have dealt with more market failures, it has become clear that the histories of the American and British markets for new doctors, and the market failures that led to their reorganization into clearinghouses, are far from unique. Other markets have failed for similar reasons, and some have been fixed in similar ways. I'll discuss common market failures we have seen in recent work on more senior medical labor markets, and also on allocation procedures that do not use prices, for school choice in New York City and Boston, and for the allocation of live-donor kidneys for transplantation. These problems were fixed by the design of appropriate clearinghouses. I will also discuss the North American labor market for new economists, in which related problems are addressed by marketplace mechanisms that leave the market relatively decentralized.

The histories of these markets suggest a number of tasks that markets and allocation systems need to accomplish to perform well. The failure to do these things causes problems that may require changes in how the marketplace is organized. I will argue that, to work well, marketplaces need to

1. provide *thickness*—that is, they need to attract a sufficient proportion of potential market participants to come together ready to transact with one another;
2. overcome the *congestion* that thickness can bring, by providing enough time, or by making transactions fast enough, so that market participants can consider enough alternative possible transactions to arrive at satisfactory ones;
3. make it *safe* to participate in the market as simply as possible
 - a. as opposed to transacting outside the marketplace, or
 - b. as opposed to engaging in strategic behavior that reduces overall welfare.

I will also remark in passing on some other lessons we have started to learn, namely that

4. some kinds of transactions are *repugnant*, and this can be an important constraint on market design.

And, on a methodological note,

5. experiments can play a role, in diagnosing and understanding market failures and successes, in testing new designs, and in communicating results to policy makers.

² Following that literature to the present would involve looking into modern designs for package auctions; see for example Cramton et al. (2006), and Milgrom (2007).

The chapter is organized as follows. The following section will describe some of the relevant history of markets for new doctors, which at different periods had to deal with each of the problems of maintaining thickness, dealing with congestion, and making it safe to participate straightforwardly in the market. In the subsequent sections I'll discuss markets in which these problems showed up in different ways.

The third section will review the recent design of regional kidney exchanges in the United States, in which the initial problem was establishing thickness, but in which problems of congestion, and, lately, making it safe for transplant centers to participate, have arisen. This is also the market most shaped by the fact that many people find some kinds of transactions repugnant. In particular, buying and selling kidneys for transplantation is illegal in most countries. So, unlike the several labor markets I discuss in this chapter, this market operates entirely without money, which will cast into clear focus how the "double coincidence of wants" problems that are most often solved with money can be addressed with computer technology (and will highlight why these problems are difficult to solve even with money, in markets like labor markets in which transactions are heterogeneous).

The fourth section will review the design of the school choice systems for New York City high schools (in which congestion was the immediate problem to be solved), and the design of the new public school choice system in Boston, in which making it safe to participate straightforwardly was the main issue. These allocation systems also operate without money.

The fifth section will discuss recent changes in the market for American gastroenterologists, who wished to adopt the kind of clearinghouse organization already in place for younger doctors, but who were confronted with some difficulties in making it safe for everyone to change simultaneously from one market organization to another. This involved making changes in the rules of the decentralized market that would precede any clearinghouse even once it was adopted.

This will bring us naturally to a discussion of changes recently made in the decentralized market for new economists in the United States.

MARKETS FOR NEW DOCTORS IN THE UNITED STATES, CANADA, AND BRITAIN³

The first job American doctors take after graduating from medical school is called a residency. These jobs are a big part of hospitals' labor force, a critical part of physicians' graduate education, and a substantial influence on their future careers. From 1900 to 1945, one way that hospitals competed for new residents was to try to hire them earlier than other hospitals. This moved the date of appointment earlier, first slowly and then

³ The history of the American medical market given here is extracted from more detailed accounts in Roth (1984, 2003, 2007).

quickly, until by 1945 residents were sometimes being hired almost two years before they would graduate from medical school and begin work.

When I studied this in Roth (1984) it was the first market in which I had seen this kind of “unraveling” of appointment dates, but today we know that unraveling is a common and costly form of market failure. What we see when we study markets in the process of unraveling is that offers not only come increasingly early, but also become dispersed in time and of increasingly short duration. So not only are decisions being made early (before uncertainty is resolved about workers’ preferences or abilities), but also quickly, with applicants having to respond to offers before they can learn what other offers might be forthcoming.⁴ Efforts to prevent unraveling are venerable; for example, Roth and Xing (1994) quote Salzman (1931) on laws in various English market from the 13th century concerning “forestalling” a market by transacting before goods could be offered in the market.⁵

In 1945, American medical schools agreed not to release information about students before a specified date. This helped control the date of the market, but a new problem emerged: hospitals found that if some of the first offers they made were rejected after a period of deliberation, the candidates to whom they wished to make their next offers had often already accepted other positions. This led hospitals to make exploding offers to which candidates had to reply immediately, before they could learn what other offers might be available, and led to a chaotic market that shortened in duration from year to year, and resulted not only in missed agreements but also in broken ones. This kind of congestion also has since been seen in other markets, and in the extreme form it took in the American medical market by the late 1940s it also constitutes a form of market failure (cf. Roth and Xing, 1997, and Avery et al., 2007, for detailed accounts of congestion in labor markets in psychology and law).

⁴ On the costs of such unraveling in some markets for which unusually good data have been available, see Niederle and Roth (2003b) on the market for gastroenterology fellows, and Fréchette et al. (2007) on the market for post-season college football bowls. For some other recent unraveled markets, see Avery et al. (2003) on college admissions; and Avery et al. (2001) on appellate court clerks. For a line of work giving theoretical insight into some possible causes of unraveling, see Li and Rosen (1998), Li and Suen (2000), Suen (2000), and Damiano et al. (2005).

⁵ “Thus at Norwich no one might forestall provisions by buying, or paying ‘earnest money’ for them before the Cathedral bell had rung for the mass of the Blessed Virgin; at Berwick-on-Tweed no one was to buy salmon between sunset and sunrise, or wool and hides except at the market-cross between 9 and 12; and at Salisbury persons bringing victuals into the city were not to sell them before broad day.” Unraveling could be in space as well as in time. Salzman also reports (p. 132) that under medieval law markets could be prevented from being established too near to an existing market, and also, for markets on rivers, nearer to the sea. “Besides injury through mere proximity, and anticipation in time, there might be damage due to interception of traffic...” Such interception was more usual in the case of waterborne traffic. In 1233 Eve de Braose complained that Richard fitz-Stephen had raised a market at Dartmouth to the injury of hers at Totnes, as ships which ought to come to Totnes were stopped at Dartmouth and paid customs there. No decision was reached, and eight years later Eve’s husband, William de Cantelupe, brought a similar suit against Richard’s son Gilbert. The latter pleaded that his market was on Wednesday and that at Totnes on Saturday; but the jury said that the market at Dartmouth was to the injury of Totnes, because Dartmouth lies between it and the sea, so that ships touched there and paid toll instead of going to Totnes; and also that cattle and sheep which used to be taken to Totnes market were now sold at Dartmouth; the market at Dartmouth was therefore disallowed.

Faced with a market that was working very badly, the various American medical associations (of hospitals, students, and schools) agreed to employ a centralized clearinghouse to coordinate the market. After students had applied to residency programs and been interviewed, instead of having hospitals make individual offers to which students had to respond immediately, students and residency programs would instead be invited to submit rank order lists to indicate their preferences. That is, hospitals (residency programs) would rank the students they had interviewed, students would rank the hospitals (residency programs) at which they had been interviewed, and a centralized clearinghouse—a matching mechanism—would be employed to produce a matching from the preference lists. Today this centralized clearinghouse is called the National Resident Matching Program (NRMP).

Roth (1984) showed that the algorithm adopted in 1952 produced a matching of students to residency programs that is *stable* in the sense defined by Gale and Shapley (1962), namely that, in terms of the submitted rank order lists, there was never a student and a residency program that were not matched to each other but would have mutually preferred to have been matched to each other than to (one of) their assigned match(es). However, changes in the market over the years made this more challenging.

For example, one change in the market had to do with the growing number of married couples graduating from American medical schools and wishing to be matched to jobs in the same vicinity. This hadn't been a problem in the 1950s, when virtually all medical students were men. Similarly, the changing nature of medical specialization sometimes produced situations in which a student needed to be simultaneously matched to two positions. Roth (1984) showed that these kinds of changes can sometimes make it impossible to find a stable matching, and, indeed, an early attempt to deal with couples in a way that did not result in a stable matching had made it difficult to attract high levels of participation by couples in the clearinghouse.

In 1995, I was invited to direct the redesign of the medical match, in response to a crisis in confidence that had developed regarding its ability to continue to serve the medical market, and whether it appropriately served student interests. A critical question was to what extent the stability of the outcome was important to the success of the clearinghouse. Some of the evidence came from the experience of British medical markets. Roth (1990, 1991b) had studied the clearinghouses that had been tried in the various regions of the British National Health Service (NHS) after those markets unraveled in the 1960s. A Royal Commission had recommended that clearinghouses be established on the American model, but since the American medical literature didn't describe in detail how the clearinghouse worked, each region of the NHS adopted a different algorithm for turning rank order lists into matches, and the unstable mechanisms had largely failed and been abandoned, while the stable mechanisms succeeded and survived.⁶

⁶ The effects of instability were different in Britain than in the US, because positions in Britain were assigned by the National Health Service, and so students were not in a position to receive other offers (and decline the positions they were matched to) as they were in the US. Instead, in Britain, students and potential employers acted in advance of unstable clearinghouses. For example, Roth

Of course, there are other differences between regions of the British NHS than how they organized their medical clearinghouses, so there was also room for controlled experiments in the laboratory on the effects of stable and unstable clearinghouses. Kagel and Roth (2000) report a laboratory experiment that compared the stable clearinghouse adopted in Edinburgh with the unstable one adopted in Newcastle, and showed that, holding all else constant, the difference in how the two clearinghouses were organized was sufficient to account for the success of the Edinburgh clearinghouse and the failure of the unstable one in Newcastle.

Roth and Peranson (1999) report on the new clearinghouse algorithm that we designed for the NRMP, which aims to always produce a stable matching. It does so in a way that makes it safe for students and hospitals to reveal their preferences.⁷ The new algorithm has been used by the NRMP since 1998, and has subsequently been adopted by over three dozen labor market clearinghouses. The empirical evidence that has developed in use is that the set of stable matchings is very seldom empty.

An interesting historical note is that the use of stable clearinghouses has been explicitly recognized as part of a pro-competitive market mechanism in American law. This came about because in 2002, sixteen law firms representing three former medical residents brought a class-action antitrust suit challenging the use of the matching system for medical residents. The theory of the suit was that the matching system was a conspiracy to hold down wages for residents and fellows, in violation of the Sherman Antitrust Act. Niederle and Roth (2003a) observed that, empirically, the wages of medical specialties with and without centralized matching in fact do not differ.⁸ The case was dismissed after the US Congress passed new legislation in 2004 (contained in Public Law 108-218)

(1991) reports that in Newcastle and Birmingham it became common for students and consultants (employers) to reach agreement in advance of the match, and then submit only each other's name on their rank order lists.

⁷ Abstracting somewhat from the complexities of the actual market, the Roth-Peranson algorithm is a modified student-proposing deferred acceptance algorithm (Gale and Shapley, 1962; see also Roth, 2008b). In simple markets, this makes it a dominant strategy for students to state their true preferences (see Roth, 1982a, 1985; Roth and Sotomayor, 1990). Although it cannot be made a dominant strategy for residency programs to state their true preferences (Roth, 1985; Sönmez, 1997), the fact that the medical market is large turns out to make it very unlikely that residency programs can do any better than to state their true preferences. This was shown empirically in Roth and Peranson (1999), and has more recently been explained theoretically by Immorlica and Mahdian (2005) and Kojima and Pathak (2009).

⁸ Bulow and Levin (2006) sketch a simple model of one-to-one matching in which a centralized clearinghouse, by enforcing impersonal wages (i.e. the same wage for any successful applicant), could cause downward pressure on wages (see also Kamecke, 1998). Subsequent analysis suggests more skepticism about any downward wage effects in actual medical labor markets. See, for example, Kojima (2007), who shows that the Bulow-Levin results don't follow in a model in which hospitals can employ more than one worker, and Niederle (2007), who shows that the results don't follow in a model that includes some of the options that the medical match actually offers patients. Crawford (2008) considers how the deferred acceptance algorithm of Kelso and Crawford (1982) could be adapted to adjust personal wages in a centralized clearinghouse (see also Artemov, 2008).

noting that the medical match is a pro-competitive market mechanism, not a conspiracy in restraint of trade. This reflected modern research on the market failures that preceded the adoption of the first medical clearinghouse in the 1950s, which brings us back to the main subject of the present chapter.⁹

To summarize, the study and design of a range of clearinghouses in the 1980s and 1990s made it clear that producing a stable matching is an important contributor to the success of a labor clearinghouse. For the purposes of the present chapter, note that such a clearinghouse can persistently attract the participation of a high proportion of the potential participants, and when it does so it solves the problem of establishing a thick market. A computerized clearinghouse like those in use for medical labor markets also solves the congestion problem, since all the operations of the clearinghouse can be conducted essentially simultaneously, in that the outcome is determined only after the clearinghouse has cleared the market. And, as mentioned briefly, these clearinghouses can be designed to make it safe for participants to reveal their true preferences, without running a risk that by doing so they will receive a worse outcome than if they had behaved strategically and stated some other preferences.

In the following sections, we'll see more about how the failure to perform these tasks can cause markets to fail.

⁹ See Roth (2003). The law states in part: "Congress makes the following findings: For over 50 years, most United States medical school seniors and the large majority of graduate medical education programs (popularly known as 'residency programs') have chosen to use a matching program to match medical students with residency programs to which they have applied.... Before such matching programs were instituted, medical students often felt pressure, at an unreasonably early stage of their medical education, to seek admission to, and accept offers from, residency programs. As a result, medical students often made binding commitments before they were in a position to make an informed decision about a medical specialty or a residency program and before residency programs could make an informed assessment of students' qualifications. This situation was inefficient, chaotic, and unfair and it often led to placements that did not serve the interests of either medical students or residency programs. The original matching program, now operated by the independent non-profit National Resident Matching Program and popularly known as 'the Match', was developed and implemented more than 50 years ago in response to widespread student complaints about the prior process.... The Match uses a computerized mathematical algorithm... to analyze the preferences of students and residency programs and match students with their highest preferences from among the available positions in residency programs that listed them. Students thus obtain a residency position in the most highly ranked program on their list that has ranked them sufficiently high among its preferences.... Antitrust lawsuits challenging the matching process, regardless of their merit or lack thereof, have the potential to undermine this highly efficient, pro-competitive, and long-standing process. The costs of defending such litigation would divert the scarce resources of our country's teaching hospitals and medical schools from their crucial missions of patient care, physician training, and medical research. In addition, such costs may lead to abandonment of the matching process, which has effectively served the interests of medical students, teaching hospitals, and patients for over half a century.... It is the purpose of this section to confirm that the antitrust laws do not prohibit sponsoring, conducting, or participating in a graduate medical education residency matching program, or agreeing to do so; and ensure that those who sponsor, conduct or participate in such matching programs are not subjected to the burden and expense of defending against litigation that challenges such matching programs under the antitrust laws."

KIDNEY EXCHANGE

Kidney transplantation is the treatment of choice for end-stage renal disease, but there is a grave shortage of transplantable kidneys. In the United States there are over 70,000 patients on the waiting list for cadaver kidneys, but in 2006 fewer than 11,000 transplants of cadaver kidneys were performed. In the same year, around 5,000 patients either died while on the waiting list or were removed from the list as “Too Sick to Transplant.” This situation is far from unique to the United States: In the UK at the end of 2006 there were over 6,000 people on the waiting list for cadaver kidneys, and only 1,240 such transplants were performed that year.¹⁰

Because healthy people have two kidneys, and can remain healthy with just one, it is also possible for a healthy person to donate a kidney, and a live-donor kidney has a greater chance of long-term success than does one from a deceased donor. However, good health and goodwill are not sufficient for a donor to be able to give a kidney to a particular patient: the patient and donor may be biologically incompatible because of blood type, or because the patient’s immune system has already produced antibodies to some of the donor’s proteins. In the United States in 2006 there were 6,428 transplants of kidneys from living donors (in the UK there were 590).

The total supply of transplantable kidneys (from deceased and living donors) clearly falls far short of the demand. But it is illegal in almost all countries to buy or sell kidneys for transplantation. This legislation is the expression of the fact that many people find the prospect of such a monetized market highly repugnant (see Roth, 2007).

So, while a number of economists have devoted themselves to the task of repealing or relaxing laws against compensating organ donors (see e.g. Becker and Elias, 2007, and the discussion of Elias and Roth, 2007), another task that faces a market designer is how to increase the number of transplants subject to existing constraints, including those that forbid monetary incentives.

It turns out that, prior to 2004, in just a very few cases, incompatible patient–donor pairs and their surgeons had managed to arrange an *exchange* of donor kidneys (sometimes called “paired donation”), when the patient in each of two incompatible patient–donor pairs was compatible with the donor in the other pair, so that each patient received a kidney from the other’s donor. Sometimes a different kind of exchange had also been accomplished, called a *list exchange*, in which a patient’s incompatible donor donated a kidney to someone who (by virtue of waiting a long time) had high priority on the waiting list for a cadaver kidney, and in return the donor’s intended patient received high priority to receive the next compatible cadaver kidney that became available. Prior

¹⁰ For US data see <<http://www.optn.org/data>> (accessed August 13, 2007; website since moved to <<http://optn.transplant.hrsa.gov>>). For UK data, see <http://www.uktransplant.org.uk/ukt/statistics/calendar_year_statistics/pdf/yearly_statistics_2006.pdf> (accessed August 13, 2007). As I update this in 2012, the number of US patients waiting for cadaver kidneys has risen to over 90,000, while in 2011 there were just barely over 11,000 transplants from cadaver kidneys (so the waiting list has grown considerably while the number of deceased donors has not).

to December 2004 only five exchanges had been accomplished at the fourteen transplant centers in New England. Some exchanges had also been accomplished at Johns Hopkins in Baltimore, and among transplant centers in Ohio. So, these forms of exchange were feasible and non-repugnant.¹¹ Why had so very few happened?

One big reason had to do with the (lack of) thickness of the market, i.e. the size of the pool of incompatible patient–donor pairs who might be candidates for exchange. When a kidney patient brought a potential donor to his or her doctor to be tested for compatibility, donors who were found to be incompatible with their patient were mostly just sent home. They were not patients themselves, and often no medical record at all was retained to indicate that they might be available. And in any event, medical privacy laws made these potential donors' medical information unavailable.

Roth et al. (2004a) showed that, in principle, a substantial increase in the number of transplants could be anticipated from an appropriately designed clearinghouse that assembled a database of incompatible patient–donor pairs. That paper considered exchanges with no restrictions on their size, and allowed list exchange to be integrated with exchange among incompatible patient–donor pairs. That is, exchanges could be a cycle of incompatible patient–donor pairs of any size such that the donor in the first pair donated a kidney to the patient in the second, the second pair donated to the third, and so on, until the cycle closed with the last pair donating to the first. And pairs that would have been interested in a list exchange in which they donated a kidney in exchange for high priority on the cadaver waiting list could be integrated with the exchange pool by having them donate to another incompatible pair in a chain that would end with donation to the waiting list.

We sent copies of that paper to many kidney surgeons, and one of them, Frank Delmonico (the medical director of the New England Organ Bank), came to lunch to pursue the conversation. Out of that conversation, which grew to include many others (and led to modifications of our original proposals), came the New England Program for Kidney Exchange, which unites the fourteen kidney transplant centers in New England to allow incompatible patient–donor pairs from anywhere in the region to find exchanges with other such pairs.

For incentive and other reasons, all such exchanges have been done simultaneously, to avoid the possibility of a donor becoming unwilling or unable to donate a kidney after that donor's intended patient has already received a kidney from another patient's donor. So, one form that congestion takes in organizing kidney exchanges is that multiple operating rooms and surgical teams have to be assembled. (A simultaneous exchange between two pairs requires four operating rooms and surgical teams, two for the nephrectomies that remove the donor kidneys, and two for the transplantations that immediately follow. An exchange involving three pairs involves six operating rooms and teams, etc.) Roth et al. (2004a) noted that large exchanges would arise relatively infrequently, but could pose logistical difficulties.

¹¹ See Rapoport (1986), Ross et al. (1997), Ross and Woodle (2000), for some early discussion of the possibility of kidney exchange, and Delmonico (2004), and Montgomery et al. (2005) for some early reports of successful exchanges.

These logistical difficulties loomed large in our early discussions with surgeons, and out of those discussions came the analysis in Roth et al. (2005a) of how kidney exchanges might be organized if only two-way exchanges were feasible. The problem of two-way exchanges can be modeled as a classic problem in graph theory, and, subject to the constraint that exchanges involve no more than two pairs, efficient outcomes with good incentive properties can be found in computationally efficient ways. When the New England Program for Kidney Exchange was founded in 2004 (Roth et al., 2005b), it used the matching software that had had been developed to run the simulations in Roth et al. (2005a,b), and it initially attempted only two-way matches (while keeping track of the potential three-way matches that were missed). This was also the case when Sönmez, Ünver and I started running matches for the Ohio-based consortium of transplant centers that eventually became the Alliance for Paired Donation.¹²

However, some transplants are lost that could have been accomplished if three-way exchanges were available. In Saidman et al. (2006) and in Roth et al. (2007), we showed that to get close to the efficient number of transplants, the infrastructure to perform both two-way and three-way exchanges would have to be developed, but that once the population of available patient–donor pairs was large enough, few transplants would be missed if exchanges among more than three pairs remained difficult to accomplish. Both the New England Program for Kidney Exchange and the Alliance for Paired Donation have since taken steps to be able to accommodate three-way as well as two-way exchanges. Being able to deal with the (six operating room) congestion required to accomplish three-way exchanges has the effect of making the market thicker, since it creates more exchange possibilities.

As noted above, another way to make the market thicker is to integrate exchange between pairs with list exchange, so that exchange chains can be considered, as well as cycles. This applies as well to how the growing numbers of non-directed (altruistic) donors are used. A non-directed (ND) donor is someone who wishes to donate a kidney without having a particular patient in mind (and whose donor kidney therefore does not require another donor kidney in exchange). The traditional way to utilize such ND donors was to have them donate to someone on the cadaver waiting list. But as exchanges have started to operate, it has now become practical to have the ND donor donate to some pair that is willing to exchange a kidney, and have that pair donate to someone on the cadaver waiting list. Roth et al. (2006) report on how and why such exchanges are now done in New England. As in traditional exchange, all surgeries are conducted simultaneously, so there are logistical limits on how long a chain is feasible.

¹² The New England Program for Kidney Exchange has since integrated our software into theirs, and conducts its own matches. The Alliance for Paired Donation originally used our software, and as the size of the exchange pool grew, the integer programming algorithms were written in software that can handle much larger numbers of pairs (Abraham et al., 2007). The papers by Roth et al. (2005a,b) were also widely distributed to transplant centers (as working papers in 2004). The active transplant program at Johns Hopkins has also begun to use software similar in design to that in Roth et al. (2004b, 2005a) to optimize pairwise matches (see Segev et al., 2005).

But we noted that, when a chain is initiated by a ND donor, it might be possible to relax the constraints that all parts of the exchange be simultaneous, since

If something goes wrong in subsequent transplants and the whole ND-chain cannot be completed, the worst outcome will be no donated kidney being sent to the waitlist and the ND donation would entirely benefit the KPD [kidney exchange] pool. (Roth et al., 2006, p. 2704)

That is, if a conventional exchange were done in a non-simultaneous way, and if the exchange broke down after some patient–donor pair had donated a kidney but before they had received one, then that pair would not only have lost the promised transplant, but also have lost a healthy kidney. In particular, the patient would no longer be in position to exchange with other incompatible patient–donor pairs. But in a chain that begins with a ND donor, if the exchange breaks down before the donation to some patient–donor pair has been made (because the previous donor in the chain becomes unwilling or unable to donate), then the pair loses the promised transplant, but is no worse off than they were before the exchange was planned, and in particular they can still exchange with other pairs in the future. So, while a non-simultaneous ND chain of donations could create an incentive to break the chain, the costs of a breach would be less than in a pure exchange, and so the benefits (in terms of longer chains) are worth exploring. The first such non-simultaneous “never ending” altruistic donor (NEAD) chain was begun by the Alliance for Paired Donation in July 2007. A week after the first patient was transplanted from an altruistic (ND) donor, her husband donated a kidney to another patient, whose mother later donated her kidney to a third patient, whose daughter donated (simultaneously) to a fourth patient, whose sister is, as I write, now waiting to donate to another patient whose incompatible donor will be willing to “pass it forward” (Rees et al., 2009a).¹³

To summarize the progress to date, the big problem facing kidney exchange prior to 2004 was the lack of thickness in the market, so that incompatible patient–donor pairs were left in the difficult search for what Jevons (1876) famously described as a double coincidence of wants (Roth et al., 2007). By building a database of incompatible patient–donor pairs and their relevant medical data, it became possible to arrange more transplants, using a clearinghouse to maximize the number (or to achieve some quality- or priority-adjusted number) of transplants subject to various constraints. The state of the art now involves both two-way and three-way cyclical exchanges and a variety of chains, either ending with a donation to someone on the cadaver waiting list or beginning with an altruistic ND donor, or both. While large simultaneous exchanges remain logistically infeasible, the fact that almost all efficient exchanges can be accomplished in cycles of no more than three pairs, together with clearinghouse technology that can efficiently

¹³ Increasing the number of patients who benefit from the altruism of a ND donor may also increase the willingness of such donors to come forward. After publicity of the first NEAD chain on *ABC World News Tonight*, July 26, 2007 (see <<http://utoledo.edu/utcommcenter/kidney>>), the Alliance for Paired Donation has had over 100 registrations on its website of people who are offering to be altruistic living ND donors (Rees, personal communication).

find such sets of exchanges, substantially reduces the problem of congestion in carrying out exchanges. And, for chains that begin with ND donors, the early evidence is that some relaxation of the incentive constraint that all surgeries be simultaneous seems to be possible.¹⁴

There remain some challenges to further advancing kidney exchange that are also related to thickness, congestion, and incentives.

Some patients have many antibodies, so that they will need very many possible donors to find one who is compatible. For that reason and others, it is unlikely that purely regional exchanges, such as presently exist, will provide adequate thickness for all the gains from exchange to be realized. Legislation has recently been passed in the US House and Senate to remove a potential legal obstacle to a national kidney exchange.¹⁵ Aside from expanding kidney exchange to national scale, another way to increase the thickness of the market would be to make kidney exchange available not just to incompatible patient–donor pairs, but also to those who are compatible but might nevertheless benefit from exchange.¹⁶

While some of the congestion in terms of actually conducting transplants has been addressed, there is still congestion associated with the time it takes to test for immunological incompatibility between patients and donors who (based on available tests) are matched to be part of an exchange. That is, antibody production can vary over time, and so a patient and donor who appear to be compatible in the database may not in fact be. Because it now sometimes takes weeks to establish this, during which time other exchanges may go forward, some exchanges are missed that could have been accomplished if the tests for compatibility were done more quickly, so that the overall pattern of exchanges could have been adjusted.

And as regional exchanges have grown to include multiple transplant centers, a new issue has come to the fore concerning how kidney exchange should be organized to give transplant centers the incentive to inform the central exchange of all of their incompatible patient–donor pairs. Consider a situation in which transplant center A has two pairs who are mutually compatible, so that it could perform an in-house exchange between

¹⁴ The Postscript describes how non-simultaneous chains have indeed come to play a very large role in kidney exchange.

¹⁵ The proposed bill (HR 710, introduced on January 29, 2007 and passed in the House on March 7, 2007, and S 487, introduced on February 1, 2007 and passed in the Senate February 15, 2007) is “To amend the National Organ Transplant Act to clarify that kidney paired donations shall not be considered to involve the transfer of a human organ for valuable consideration.” Kidney exchange is also being organized in the UK; see <[http://www.uktransplant.org.uk/ukt/about_transplants/organ_allocation/kidney_\(renal\)/living_donation/paired_donation_matching_scheme.jsp](http://www.uktransplant.org.uk/ukt/about_transplants/organ_allocation/kidney_(renal)/living_donation/paired_donation_matching_scheme.jsp)>. The first British exchange was carried out on July 4, 2007 (see the BBC report at <<http://news.bbc.co.uk/1/hi/health/7025448.stm>>).

¹⁶ For example, a compatible middle-aged patient–donor pair, and an incompatible patient–donor pair in which the donor is a twenty-five-year-old athlete could both benefit from exchange. Aside from increasing the number of pairs available for exchange, this would also relieve the present shortage of donors with blood type O in the kidney exchange pool, caused by the fact that O donors are only rarely incompatible with their intended recipient. Simulations on the robust effects of adding compatible patient–donor pairs to the exchange pool are found in Roth et al. (2004a, 2005b), and in Gentry et al. (2007).




FIGURE 1.1. Potential kidney exchanges between patient–donor pairs at multiple centers. Double-headed arrows indicate that the connected pairs are compatible for exchange, i.e. the patient in one pair is compatible with the donor in the other. Pairs A1 and A2 are both from transplant center A; pairs B and C are from different transplant centers. Transplant center A, which sees only its own pairs, can conduct an exchange among its pairs A1 and A2 since they are compatible, and, if it does so, this will be the only exchange, resulting in two transplants. However, if in Figure 1a transplant center A makes its pairs available for exchange with other centers, then the exchanges will be A1 with B and A2 with C, resulting in four transplants. However, in Figure 1b the suggested exchange might be A1 with B, which would leave the patient in A2 without a transplant. Faced with this possibility (and not knowing if the situation is as in 1a or 1b) transplant center A might choose to transplant A1 and A2 by itself, without informing the central exchange.

these two pairs. If the mutual compatibilities are as shown in Figure 1.1a, then if these two pairs exchange with each other, only those two transplants will be accomplished. If instead the pairs from transplant center A were matched with the pairs from the other centers, as shown in Figure 1.1a, four transplants could be accomplished (via exchanges of pair A1 with pair B, and pair A2 with C).

But, note that if the situation had been that of Figure 1.1b, then transplant center A runs the risk that if it informs the central exchange of its pairs, then the recommended exchange will be between A1 and B, since B has high priority (e.g. B is a child). This would mean that pair A2 did not get a kidney, as they would have if A1 and A2 had exchanged in-house. So, the situation facing transplant center A, not knowing what pairs will be put forward for exchange by the other transplant centers, is that it can assure itself of doing two transplants for its patients in pairs A1 and A2, but it is not guaranteed two transplants if it makes the pairs available for exchange and the situation is as in Figure 1.1b. If this causes transplant centers to withhold those pairs they can transplant by themselves, then a loss to society results where the situation is as in Figure 1.1a. (In fact, if transplant centers withhold those pairs they can exchange in-house, then primarily hard-to-match pairs will be offered for exchange, and the loss will be considerable.)

One remedy is to organize the kidney exchange clearinghouse in a way that guarantees center A that any pairs it could exchange in-house will receive transplants. This

would allow the maximal number of transplants to be achieved in the situation depicted in Figure 1.1a, and it would mean that in the situation depicted in Figure 1.1b the exchange between A1 and A2 would be made (and so the high-priority pair B would not participate in exchange, just as they would not have if pairs A1 and A2 had not been put forward). This is a bit of a hard discussion to have with surgeons, who find it repugnant that, for example, the child patient in pair B would receive lower priority than pairs A1 and A2 just because of the accident that they were mutually compatible and were being treated at the same transplant center. (Needless to say, if transplant center A withholds its pairs and transplants them in-house, they effectively have higher priority than pair B, even if no central decision to that effect has been made.) But this is an issue that will have to be resolved, because the full participation of all transplant centers substantially increases the efficiency of exchange.

Note that, despite all the detailed technical particulars that surround the establishment of kidney exchange programs, and despite the absence of money in the kidney exchange market, we can recognize some of the basic lessons of market design that were also present in designing labor market clearinghouses. The first issue was making the market thick, by establishing a database of patient–donor pairs available to participate in exchange. Then issues of congestion had to be dealt with, so that the clearinghouse could identify exchanges involving sufficiently few pairs (initially two, now three) for transplants to be done simultaneously. Simultaneity is related to making sure that everyone involved in an exchange never has an incentive not to go forward with it, but as exchanges have grown to include multiple transplant centers, there are also incentive issues to be resolved in making it safe for a transplant center to enroll all of its eligible pairs in the central exchange.

SCHOOL CHOICE

Another important class of allocation problems in which no money changes hands is the assignment of children to big-city public schools, based both on the preferences of students and their families, and on the preferences of schools, or on city priorities. Because public school students must use whatever system local authorities establish, establishing a thick market is not the main problem facing such systems. (Although how well a school choice system works may influence how many children ultimately attend city schools.) But how well a school choice system works still has to do with how effectively it deals with congestion, and how safe it makes it for families to straightforwardly reveal their preferences.

My colleagues and I were invited to help design the current New York City (NYC) high-school choice program, chiefly because of problems the old decentralized system had in dealing with congestion. In Boston we were invited to help design the current school choice system because the old system, which was itself a centralized

clearinghouse, did not make it safe for families to state their preferences.¹⁷ In both Boston and NYC the newly designed systems incorporate clearinghouses to which students (and, in NYC, schools) submit preferences. Although another alternative was considered in Boston, both Boston and NYC adopted clearinghouses similar to the kinds of stable clearinghouses used in medical labor markets (powered by a student-proposing deferred acceptance algorithm), adapted to the local situations. For my purpose in the present chapter, I'll skip any detailed discussion of the clearinghouse designs, except to note that they make it safe for students and families to submit their true preferences. Instead, I'll describe briefly what made the prior school choice systems congested or risky.¹⁸

In NYC, well over 90,000 students a year must be assigned to over 500 high-school programs. Under the old system, students were asked to fill out a rank order list of up to five programs. These lists were then copied and sent to the schools. Subject to various constraints, schools could decide which of their applicants to accept, waitlist, or reject. Each applicant received a letter from the NYC Department of Education with the decisions of the schools to which she or he had applied, and applicants were allowed to accept no more than one offer, and one waitlist. This process was repeated: after the responses to the first letter were received, schools with vacant positions could make new offers, and after replies to the second letter were received, a third letter with new offers was sent. Students not assigned after the third step were assigned to their zoned schools, or assigned via an administrative process. There was an appeals process, and an "over the counter" process for assigning students who had changed addresses, or were otherwise unassigned before school began.

Three rounds of processing applications to no more than five out of more than 500 programs by almost 100,000 students was insufficient to allocate all the students. That is, this process suffered from congestion (in precisely the sense explored in Roth and Xing, 1997): not enough offers and acceptances could be made to clear the market. Only about 50,000 students received offers initially, about 17,000 of whom received multiple offers. And when the process concluded, approximately 30,000 students had been assigned to a school that was nowhere on their choice list.

Three features of this process particularly motivated NYC Department of Education's desire for a new matching system. First were the approximately 30,000 students not assigned to a school they had chosen. Second, students and their families had to be strategic in their choices. Students who had a substantial chance of being rejected by their true first-choice school had to think about the risk of listing it first, since, if one of their lower-choice schools took students' rankings into account in deciding on admissions, they might have done better to list it first. (More on this in a

¹⁷ The invitation to meet with Boston Public Schools came after a newspaper story recounted the difficulties with the Boston system, as described in Abdulkadiroğlu and Sönmez (2003). For subsequent explorations of the old Boston system, see Chen and Sönmez (2006), Ergin and Sönmez (2006), Pathak and Sönmez (2008), and Abdulkadiroğlu et al. (2007).

¹⁸ The description of the situation in NYC is from Abdulkadiroğlu et al. (2005a); for Boston see Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu et al. (2005b, 2007).

moment, in the discussion of Boston schools.) Finally, the many unmatched students, plus those who may not have indicated their true preferences (and the consequent instability of the resulting matching) gave schools an incentive to be strategic: a substantial number of schools managed to conceal capacity from the central administration, thus preserving places that could be filled later with students unhappy with their assignments.

As soon as NYC adopted a stable clearinghouse for high-school matching (in 2003, for students entering high school in 2004), the congestion problem was solved; only about 3,000 students a year have had to be assigned administratively since then, down from 30,000 (and many of these are students who for one reason or another fail to submit preference lists). In addition, in the first three years of operation, schools learned that it was no longer profitable to withhold capacity, and the resulting increase in the availability of places in desirable schools resulted in a larger number of students receiving their first choices, second choices, and so forth from year to year. Finally, as submitted rank order lists have begun to more reliably reflect true preferences, these have begun to be used as data for the politically complex process of closing or reforming undesirable schools (Abdulkadiroğlu et al., 2005a, 2009).

In Boston, the problem was different. The old school choice system there made it risky for parents to indicate their true first-choice school if it was not their local school. The old system was simple in conception: parents ranked schools, and the algorithm tried to give as many families as possible their first-choice school. Where the capacity of a school was less than the number of students who ranked it first, priority was given to students who had siblings in the school, or who lived within walking distance, or, finally, who had been assigned a good lottery number. After these assignments were made, the algorithm tried to match as many remaining students as possible with their second-choice school, and so on. The difficulty facing families was that, if they ranked a popular school first and weren't assigned to it, they might find that by the time they were considered for their second-choice school, it was already filled with people who had ranked it first. So, a family who had a high priority for their second-choice school (e.g. because they lived close to it), and could have been assigned to it if they had ranked it first, might no longer be able to get in if they ranked it second.

As a consequence, many families were faced with difficult strategic decisions, and some families devoted considerable resources to gathering relevant information about the capacities of schools, how many siblings would be enrolling in kindergarten, etc. Other families were oblivious to the strategic difficulties, and sometimes suffered the consequences; if they listed popular schools for which they had low priority, they were often assigned to schools they liked very little.

In Boston, the individual schools are not actors in the school choice process, and so there was a wider variety of mechanisms to choose from than in New York. My colleagues and I recommended two possibilities that were *strategy-proof* (in the sense that they make it a dominant strategy for students and families to submit their true preferences), and which thus would make it safe for students to submit their true preferences

(Abdulkadiroğlu et al., 2005b, 2007).¹⁹ This proved to be decisive in persuading the Boston School Committee to adopt a new algorithm. Then Superintendent of Schools, Thomas Payzant, wrote, in a 2005 memo to the School Committee:

The most compelling argument for moving to a new algorithm is to enable families to list their true choices of schools without jeopardizing their chances of being assigned to any school by doing so.

Superintendent Payzant further wrote:

A strategy-proof algorithm levels the playing field by diminishing the harm done to parents who do not strategize or do not strategize well.

Making the school choice system safe to participate in was critical in the decision of Boston public schools to move from a clearinghouse that was not strategy-proof to one that was. Different issues of safety were critical in the market for gastroenterologists, discussed next.

GASTROENTEROLOGISTS²⁰

An American medical graduate who wishes to become a gastroenterologist first completes three years of residency in internal medicine, and then applies for a job as a fellow in gastroenterology, a subspecialty of internal medicine.²¹ The market for gastroenterology fellows was organized via a stable labor market clearinghouse (a “match”) from 1986 through the late 1990s, after which the match was abandoned (following an unexpected shock to the supply and demand for positions in 1996; see McKinney et al., 2005). This provided an opportunity to observe the unraveling of a market as it took place. From the late 1990s until 2006, offers of positions were made increasingly far in advance of employment (moving back to almost two years in advance, so that candidates were often being interviewed early in their second year of residency). Offers also became dispersed in time, and short in duration, so that candidates faced a thin market. One consequence was that the market became much more local than it had been, with gastroenterology fellows more likely to be recruited at the same hospital at which they had worked as a resident (Niederle and Roth, 2003b; Niederle et al., 2006).

Faced with these problems, the various professional organizations involved in the market for gastroenterology fellows agreed to try to resume using a centralized

¹⁹ In addition to the student-proposing deferred acceptance algorithm that was ultimately adopted, we proposed a variation of the “top trading cycles” algorithm originally explored by Shapley and Scarf (1974), which was shown to be strategy-proof by Roth (1982b), and which was extended, and explored in a school choice context, by Abdulkadiroğlu and Sönmez (1999, 2003).

²⁰ A much more thorough treatment of the material in this section is given in Niederle and Roth (2009b).

²¹ The American system of residents and fellows is similar but not precisely parallel to the system in the UK of house officers and registrars, which has also recently faced some problems of market design.

clearinghouse, to be operated one year in advance of employment. However, this raised the question of how to make it safe for program directors and applicants to wait for the clearinghouse, which would operate almost a year later than hiring had been accomplished in the immediate past. Program directors who wanted to wait for the match worried that if their competitors made early offers, then applicants would lose confidence that the match would work and consequently would accept those early offers. That is, in the first year of a match, applicants might not yet feel safe to reject an early offer in order to wait for the match. Program directors who worried about their competitors might thus be more inclined to make early offers themselves.

The gastroenterology organizations did not feel able to directly influence the hiring behavior of programs that might not wish to wait for the match. Consequently we recommended that policies be adopted that would allow applicants who wished to wait for the match to more effectively deal with early offers themselves (Niederle et al., 2006). We modeled our recommendation on the policies in place in the American market for graduate school admission. In this market, a policy (adopted by the large majority of universities) states that offers of admission and financial support to graduate students should remain open until April 15.

Students are under no obligation to respond to offers of financial support prior to April 15; earlier deadlines for acceptance of such offers violate the intent of this Resolution. In those instances in which a student accepts an offer before April 15, and subsequently desires to withdraw that acceptance, the student may submit in writing a resignation of the appointment at any time through April 15. However, an acceptance given or left in force after April 15 commits the student not to accept another offer without first obtaining a written release from the institution to which a commitment has been made. Similarly, an offer by an institution after April 15 is conditional on presentation by the student of the written release from any previously accepted offer. It is further agreed by the institutions and organizations subscribing to the above Resolution that a copy of this Resolution should accompany every scholarship, fellowship, traineeship, and assistantship offer.” (See <<http://www.cgsnet.org/april-15-resolution>>)

This of course makes early exploding offers much less profitable. A program that might be inclined to insist on an against-the-rules early response is discouraged from doing so, because they can’t “lock up” a student to whom they make such an offer, because accepting such an offer does not prevent the student from later receiving and accepting a preferred offer.²²

A modified version of this policy was adopted by all four major gastroenterology professional organizations, the American Gastroenterological Association (AGA), the American College of Gastroenterology (ACG), the American Society for

²² Niederle and Roth (2009a) study in the laboratory the impact of the rules that govern the types of offers that can be made (with or without a very short deadline) and whether applicants can change their minds after accepting an early offer. In the uncongested laboratory environments we studied, eliminating the possibility of making exploding offers, or making early acceptances non-binding, prevents the markets from operating inefficiently early.

Gastrointestinal Endoscopy (ASGE), and the American Association for the Study of Liver Diseases (AASLD), regarding offers made before the (new) match. The resolution states, in part:

The general spirit of this resolution is that each applicant should have an opportunity to consider all programs before making a decision and be able to participate in the Match....It therefore seeks to create rules that give both programs and applicants the confidence that applicants and positions will remain available to be filled through the Match and not withdrawn in advance of it. This resolution addresses the issue that some applicants may be persuaded or coerced to make commitments prior to, or outside of, the Match....Any applicant may participate in the matching process...by...resigning the accepted position if he/she wishes to submit a rank order list of programs....The spirit of this resolution is to make it unprofitable for program directors to press applicants to accept early offers, and to give applicants an opportunity to consider all offers....

The gastroenterology match for 2007 fellows was held on June 21, 2006, and succeeded in attracting 121 of the 154 eligible fellowship programs (79%). Of the positions offered in the match, 98% were filled through the match, and so it appears that the gastroenterology community succeeded in making it safe to participate in the match, and thus in changing the timing and thickness of the market, while using a clearinghouse to avoid congestion.

The policies adopted by gastroenterologists prior to their match make clear that market design in this case consists not only of the “hardware” of a centralized clearinghouse, but also of the rules and understandings that constitute elements of “market culture.” This leads us naturally to consider how issues of timing, thickness, and congestion are addressed in a market that operates without any centralized clearinghouse.

MARKET FOR NEW ECONOMISTS

The North American market for new PhDs in economics is a fairly decentralized one, with some centralized marketplace institutions, most of them established by the American Economics Association (AEA).²³ Some of these institutions are of long standing, while others have only recently been established. Since 2005 the AEA has had an Ad Hoc Committee on the Job Market, charged with considering ways in which the market for economists might be facilitated.²⁴

²³ This is not a closed market, as economics departments outside North America also hire in this market, and as American economics departments and other employers often hire economists educated elsewhere. But a large part of the market involves new American PhDs looking for academic positions at American colleges and universities. See Cawley (2006) for a description of the market aimed at giving advice to participants, and Siegfried and Stock (2004) for some descriptive statistics.

²⁴ At the time of writing its members were Alvin E. Roth (chair), John Cawley, Philip Levine, Muriel Niederle, and John Siegfried, and the committee had received assistance from Peter Coles, Ben Greiner, and Jenna Kutz.

Roughly speaking, the main part of this market begins each year in the early fall, when economics departments advertise for positions. Positions may be advertised in many ways, but a fairly complete picture of the academic part of the market can be obtained from the AEA's monthly publication *Job Openings for Economists* (JOE), which provides a central location for employers to advertise and for job seekers to see who is hiring (<http://www.aeaweb.org/joe>). Graduate students nearing completion of their PhDs answer the ads by sending applications, which are followed by letters of reference, most typically from their faculty advisors.²⁵

Departments often receive several hundred applications (because it is easy for applicants to apply to many schools), and junior recruiting committees work through the late fall to read applications, papers, and letters, and to seek information through informal networks of colleagues, to identify small subsets of applicants they will invite for half-hour preliminary interviews at the annual AEA meeting in early January. This is part of a very large annual set of meetings, of the Allied Social Science Associations (ASSA), which consist of the AEA and almost fifty smaller associations. Departments reserve suites for interviewing candidates at the meeting hotels, and young economists in new suits commute up and down the elevators, from one interview to another, while recruiting teams interview candidates one after the other, trading off with their colleagues throughout long days. While the interviews in hotel suites are normally prearranged in December, the meetings also host a spot market, in a large hall full of tables, at which both academic and non-academic employers can arrange at the last minute to meet with candidates. The spot market is called the Illinois Skills Match (because it is organized in conjunction with the Illinois Department of Employment Security).

These meetings make the early part of the market thick, by providing an easy way for departments to quickly meet lots of candidates, and by allowing candidates to efficiently introduce themselves to many departments. This largely controls the starting time of the market.²⁶ Although a small amount of interviewing goes on beforehand, it is quite rare to hear of departments that make offers before the meetings, and even rarer to hear of departments pressing candidates for replies before the meetings.²⁷

²⁵ These applications are usually sent through the mail, but now often also via email and on webpages set up to receive them. Applicants typically apply to departments individually, by sending a letter accompanied by their curriculum vitae and job market paper(s) and followed by their letters of reference. Departments also put together "packages" of their graduating students who are on the market, consisting of curricula vitae, job market papers, and letters of reference, and these are sent by mail and/or posted on department websites (without the letters of reference). In 2007 a private organization, EconJobMarket.org, offered itself as a central repository of applications and letters of reference on the web. The European Economics Association in collaboration with the Asociación Española de Economía has initiated a similar repository at <http://jobmarketeconomist.com>.

²⁶ The situation is different in Europe, for example, where hiring is more dispersed in time. In an attempt to help create a thicker European market, the Royal Economic Society held a "PhD presentations event" for the first time in late January 2006. Felli and Sutton (2006) remark that "The issue of timing, unsurprisingly, attracted strong comment..."

²⁷ While the large-scale interviewing at the annual meetings has not been plagued by gradual unraveling, some parts of the market have broken off. In the 1950s, for example, the American Marketing Association used to conduct job market meetings at the time of the ASSA meetings, but for a

But while the preliminary interviewing part of the market is thick, it is congested. A dedicated recruiting committee might be able to interview thirty candidates, but not a hundred, and hence can meet only a small fraction of the available applicants. Thus the decision of whom to interview at the meetings is an important one, and for all but elite schools a strategic one as well. That is, while a few departments at the top of the pecking order can simply interview the candidates they like best, a lower-ranked department that uses all its interview slots to interview the same candidates who are interviewed by the elite schools is likely to find that it cannot convert its initial interviews into new faculty hires. Thus most schools have to give at least some thought not only to how much they like each candidate, but to how likely it is that they can successfully hire that candidate. This problem is only made more difficult by the fact that students can easily apply for many positions, so the act of sending an application does not itself send a strong signal of how interested the candidate might be. The problem may be particularly acute for schools in somewhat special situations, such as liberal arts colleges, or British and other non-American universities in which English is the language of instruction, since these may be concerned that some students who strongly prefer positions at North American research universities may apply to them only as insurance.

Following the January meetings, the market moves into a less organized phase, in which departments invite candidates for “flyouts,” day-long campus visits during which the candidate will make a presentation and meet a substantial portion of the department faculty and perhaps a dean. Here, too, the market is congested, and departments can fly out only a small subset of the candidates they have interviewed at the meetings, because of the costs of various sorts.²⁸ This part of the market is less well coordinated in time: some departments host flyouts in January, while others wait until later. Some departments try to complete all their flyouts before making any offers, while others make offers while still interviewing. And some departments make offers that come with moderate deadlines of two weeks or so, which may nevertheless force candidates to reply to an offer before knowing what other offers might be forthcoming.²⁹

By late March, the market starts to become thin. For example, a department that interviewed twenty people at the meetings, invited six for flyouts, made offers to two, and was rejected by both, may find that it is now difficult to assess which candidates it did not interview may still be on the market. Similarly, candidates whose interviews

long time it has held its job market in August, a year before employment will begin, with the result that assistant professors of marketing are often hired before having made as much progress on their dissertations as is the case for economists (Roth and Xing, 1994).

²⁸ These costs arise not only because budgets for airfares and hotels may be limited, but also because faculties quickly become fatigued after too many seminars and recruiting dinners.

²⁹ In 2002 and 2003 Georg Weizsacker, Muriel Niederle, Dorothea Kubler, and I conducted surveys of economics departments regarding their hiring practices, asking in particular about what kinds of deadlines, if any, they tended to give when they made offers to junior candidates. Loosely speaking, the results suggested that departments that were large, rich, and elite often did not give any deadlines (and sometimes were able to make all the offers they wanted to make in parallel, so that they would not necessarily make new offers upon receiving rejections). Less well endowed departments often gave candidates deadlines, although some were in a position to extend the deadline for candidates who seemed interested but needed more time.

and flyouts did not result in job offers may find it difficult to know which departments are still actively searching. To make the late part of the market thicker, the first thing our AEA job market committee did was to institute a “scramble” webpage through which departments with unfilled positions and applicants still on the market could identify each other (see *Guide to the Economics Job Market Scramble* at <<http://www.aeaweb.org/joe/scramble/guide.pdf>>). For simplicity, the scramble webpage was passive (i.e. it didn’t provide messaging or matching facilities): it simply announced the availability of any applicant or department who chose to register. The scramble webpage operated for the first time in the latter part of the 2005–06 job market, when it was open for registrants between March 15 and 20, and was used by 70 employers and 518 applicants (of whom only about half were new, 2006 PhDs). It was open only briefly, so that its information provided a snapshot of the late market, which didn’t have to be maintained to prevent the information from becoming stale.

The following year our committee sought to alleviate some of the congestion surrounding the selection of interview candidates at the January meetings, by introducing a signaling mechanism through which applicants could have the AEA transmit to no more than two departments a signal indicating their interest in an interview at the meetings. The idea was that, by limiting applicants to two signals, each signal would have some information value that might not be contained merely in the act of sending a department an application, and that this information might be helpful in averting coordination failures.³⁰ The signaling mechanism operated for the first time in December 2006, and about 1,000 people used it to send signals.³¹

³⁰ For a simple conceptual example of how a limited number of signals can improve welfare, consider a market with two applicants and two employers, in which there is only time for each employer to make one offer, and each applicant can take at most one position. Even if employers and applicants wish only to find a match, and have no preference with whom they match, there is a chance for signals to improve welfare by reducing the likelihood of coordination failure. In the absence of signals, there is a symmetric equilibrium in which each firm makes an offer to each worker with equal probability, and at this equilibrium, half the time one worker receives two offers, and so one worker and one employer remain unmatched. If the workers are each permitted to send one signal beforehand, and if each worker sends a signal to each firm with equal probability, then if firms adopt the strategy of making an offer to an applicant who sends them a signal, the chance of coordination failure is reduced from one-half to one-quarter. If workers have preferences over firms, the welfare gains from reducing coordination failure can be even larger. For recent treatments of signaling and coordination, see Coles et al. (forthcoming), Lee and Schwarz (2007a,b), Lien (2007), and Stack (2007). See also Abdulkadiroğlu et al. (2011), who discuss allowing applicants to influence tie-breaking by signaling their preferences in a centralized clearinghouse that uses a deferred acceptance algorithm.

³¹ The document “Signaling for Interviews in the Economics Job Market,” at <<http://www.aeaweb.org/joe/signal/signaling.pdf>> includes the following advice:

“**Advice to Departments:** Applicants can only send two signals, so if a department *doesn’t* get a signal from some applicant, that fact contains almost no information. (See advice to applicants, below, which suggests how applicants might use their signals). But because applicants can send only two signals, the signals a department *does* receive convey valuable information about the candidate’s interest.” “A department that has more applicants than it can interview can use the signals to help break ties for interview slots, for instance. Similarly, a department that receives applications from some candidates who it thinks are unlikely to really be interested (but might be submitting many applications out of excessive risk aversion) can be reassured of the candidate’s interest if the department receives one of the

Both the scramble and the signaling facility attracted many users, although it will take some time to assess their performance. Like the JOE and the January meetings, they are marketplace institutions that attempt to help the market provide thickness and deal with congestion.

DISCUSSION

In the tradition of market design, I have concentrated on the details of particular markets, from medical residents and fellows to economists, and from kidney exchange to school choice. But, despite their very different details, these markets, like others, struggle to provide thickness, to deal with the resulting congestion, and to make it safe and relatively simple to participate. While the importance of thick markets has been understood by economists for a long time, my impression is that issues of congestion, safety, and simplicity were somewhat obscured when the prototypical market was thought of as a market for a homogeneous commodity.³²

Thickness in a market has many of the properties of a public good, so it is not surprising that it may be hard to provide it efficiently, and that free riders have to be resisted, whether in modern markets with a tendency to unravel, or in medieval markets with rules against “forestalling.” Notice that providing thickness blurs the distinction between centralized and decentralized markets, since marketplaces—from traditional farmers’ markets, to the AEA job market meetings, to the New York Stock Exchange—provide thickness by bringing many participants to a central place. The possibility of having the market perform other centralized services, as clearinghouses or signaling mechanisms do, has only grown now that such central places can also be electronic, on the Internet or elsewhere. And issues of thickness become if anything more important when there are network externalities or other economies of scope.³³

candidate’s two signals. A department that receives a signal from a candidate will likely find it useful to open that candidate’s dossier and take one more look, keeping in mind that the candidate thought it worthwhile to send one of his two signals to the department.”

“Advice to Applicants: The two signals should **not** be thought of as indicating your top two choices. Instead, you should think about which two departments that you are interested in would be likely to interview you if they receive your signal, but not otherwise (see advice to departments, above). You might therefore want to send a signal to a department that you like but that might otherwise doubt whether they are likely to be able to hire you. Or, you might want to send a signal to a department that you think might be getting many applications from candidates somewhat similar to you, and a signal of your particular interest would help them to break ties. You might send your signals to departments to whom you don’t have other good ways of signaling your interest.”

³² Establishing thickness, in contrast, is a central concern even in financial markets; see for example the market design (“market microstructure”) discussions of how markets are organized at their daily openings and closings, such as Biais et al. (1999) on the opening call auction in the Paris Bourse and Kandel et al. (2007) on the closing call auctions in the Borsa Italiana and elsewhere.

³³ Thickness has received renewed attention in the context of software and other “platforms” that serve some of the functions of marketplaces, such as credit cards, which require large numbers of both

Congestion is especially a problem in markets in which transactions are heterogeneous, and offers cannot be made to the whole market. If transactions take even a short time to complete, but offers must be addressed to particular participants (as in offers of a job, or to purchase a house), then someone who makes an offer runs the risk that other opportunities may disappear while the offer is being considered. And even financial markets (in which offers can be addressed to the whole market) experience congestion on days with unusually heavy trading and large price movements, when prices may change significantly while an order is being processed, and some orders may not be able to be processed at all. As we have seen, when individual participants are faced with congestion, they may react in ways that damage other properties of the market, for example if they try to gain time by transacting before others.³⁴

Safety and simplicity may constrain some markets differently than others. Parents engaged in school choice may need more of both than, say, bidders in very-high-value auctions of the sort that allow auction experts to be hired as consultants. But even in billion-dollar spectrum auctions, there are concerns that risks to bidders may deter entry, or that unmanageable complexity in formulating bids and assessing opportunities at each stage may excessively slow the auction.³⁵ Somewhere in between, insider trading laws with criminal penalties help make financial markets safe for non-insiders to participate. And if it is risky to participate in the market, individual participants may try to manage their risk in ways that damage the market as a whole, such as when transplant centers withhold patients from exchange, or employers make exploding offers before applicants can assess the market, or otherwise try to prevent their trading counterparties from being able to receive other offers.³⁶

In closing, market design teaches us both about the details of market institutions and about the general tasks markets have to perform. Regarding details, the word “design” in “market design” is not only a verb, but also a noun, so economists can help to design some markets, and profitably study the design of others. And I have argued in this chapter that among the general tasks markets have to perform, difficulties in providing

consumers and merchants; see for example Evans and Schmalensee (1999) and Evans et al. (2006); and see Rochet and Tirole (2006), who concentrate on how the price structure for different sides of the market may be an important design feature.

³⁴ The fact that transactions take time may in some markets instead inspire participants to try to transact very late, near the market close, if that will leave other participants with too little time to react. See for example the discussion of very late bids (“sniping”) on eBay auctions in Roth and Ockenfels (2002), and Arieli et al. (2005).

³⁵ Bidder safety lies behind discussions both of the “winner’s curse” and collusion (cf. Kagel and Levin 2002; Klemperer, 2004), as well as of the “exposure problem” that faces bidders who wish to assemble a package of licenses in auctions that do not allow package bidding (see e.g. Milgrom, 2007). And simplicity of the auction format has been addressed in experiments prior to the conduct of some (U.S.) Federal Communications Commission (FCC) auctions (see e.g. Plott, 1997). Experiments have multiple uses in market design, not only for investigation of basic phenomena, and small-scale testing of new designs, but also in the considerable amount of explanation, communication, and persuasion that must take place before designs can be adopted in practice.

³⁶ For example, Roth and Xing (1994) report that in 1989 some Japanese companies scheduled recruiting meetings on the day an important civil service exam was being given, to prevent their candidates from also applying for government positions.

thickness, dealing with congestion, and making participation safe and simple are often at the root of market failures that call for new market designs.

I closed my 1991 *Economic Journal* article (quoted in the introduction) on a cautiously optimistic note that, as a profession, we would rise to the challenge of market design, and that doing so would teach us important lessons about the functioning of markets and economic institutions. I remain optimistic on both counts.

POSTSCRIPT 2012: WHAT HAVE WE LEARNED FROM MARKET DESIGN LATELY?³⁷

The design of new marketplaces raises new theoretical questions, which sometimes lead to progress in economic theory. Also, after a market has been designed, adopted, and implemented, it is useful to monitor how things are going, to find out if there are problems that still need to be addressed. In this update, I'll briefly point to developments of each of these kinds since the publication of Roth (2008a), "What have we learned from market design?" I'll again discuss theoretical results only informally, to avoid having to introduce the full apparatus of notation and technical assumptions. And while I will try to separate "theoretical" and "operational" issues for clarity, what will really become clear is how closely theoretical and operational issues are intertwined in practical market design.

In Roth (2008a) I described how marketplace design often involves attracting enough participants to make a market thick, dealing with the congestion that can result from attracting many participants, and making participation in the market safe and simple. Accomplishing these tasks requires us to consider, among other things, the strategy sets of the participants, the behavior elicited by possible market designs, and the stability of the resulting outcomes (see e.g. Roth, 2002; Roth and Sotomayor, 1990). To bring theory to bear on a practical problem, we need to create a simple model that allows these issues to be addressed. In what follows, I'll discuss how sometimes an initially useful simple model becomes less useful as the marketplace changes, or as new problems have to be addressed, and how this feeds back to modifications of the original model, and to new theory developed with the help of the new models.

School choice

Theoretical issues

School assignment systems face different problems in different cities. In NYC, high-school assignment had a strong resemblance to the problems facing labor markets for

³⁷ An earlier update, in Spanish, appeared in Roth (2011).

medical school graduates. In both cases, a large number of people have to be matched with a large number of positions at around the same time. And in both cases, the “positions” are in fact strategic players: NYC high-school principals, like directors of medical residency programs, have preferences over whom they match with, and have some strategic flexibility in meeting their goals. So it made sense to think of the NYC high-school assignment process as a two-sided matching market that needed to reach a stable matching—one in which no student and school would prefer to be matched to one another than to accept their assigned matches—in order to damp down some of the strategic behavior that made it hard for the system to work well. And in NYC, as in the medical residency match, there were compelling reasons to choose the applicant-optimal stable matching mechanism—implemented via a student-proposing deferred acceptance algorithm—that makes it safe for applicants to reveal their true preferences.

However, there is an important difference between labor markets and school choice. In a labor market like the one for medical graduates, assuming that the parties have strict preferences (and requiring the graduates to rank order them) probably doesn’t introduce much distortion into the market. But in a school choice setting, schools in many cases have (and are often required to have) very large indifference classes, i.e. very many students between whom they can’t distinguish. So the question of tie-breaking arises: when there are enough places in a given school to admit only some of a group of otherwise equivalent students, who should get the available seats?

How to do tie-breaking was one of the first questions we confronted in the design of the NYC high-school match, and we had to make some choices among ways to break ties by lottery. In particular, we considered whether to give each student a single number to be used for tie-breaking at every school (single tie-breaking), or to assign numbers to each student at each school (multiple tie-breaking). Computations with simulated and then actual submitted preferences indicated that single tie-breaking had superior welfare properties. Subsequent theoretical and empirical work has clarified the issues involved in tie-breaking. A simple example with just one-to-one matching is all that will be needed to explain, but first it will be helpful to look at how the deferred acceptance algorithm works. (For a description of how the algorithm is adapted to the complexities of the NYC school system, see Abdulkadiroğlu et al., 2009.)

The basic deferred acceptance algorithm with tie-breaking proceeds as follows:

- Step 0.0: Students and schools *privately*³⁸ submit preferences (and school preferences may have ties, i.e. schools may be indifferent between some students).

³⁸ One feature of the old NYC high-school assignment process was that schools saw how students ranked them, and quite a few schools would only admit students who had ranked them first. Of course, if in the new system schools had still been permitted to see students’ rank order lists, even a student-proposing deferred acceptance algorithm would not be strategy-proof. The proof that the student-proposing deferred acceptance algorithm makes it a dominant strategy for students to state their true preferences incorporates the assumption that preference lists are private, through the assumption that the strategy sets available to the players consist of preference lists as a function (only) of own preferences, so that schools’ strategies do not include the possibility of making their preference list contingent on the preference lists submitted by students (see Roth, 1982).

- Step 0.1: Arbitrarily break all ties in preferences.
- Step 1: Each student “applies” to her or his first choice. Each school tentatively assigns its seats to its applicants one at a time in their priority order. Any remaining applicants are rejected.
- ...
- Step k : Each student who was rejected in the previous step applies to her or his next choice if one remains. Each school considers the students it has been holding together with its new applicants and tentatively assigns its seats to these students one at a time in priority order. Any remaining applicants are rejected.
- The algorithm terminates when no student application is rejected, and each student is assigned her or his final tentative assignment.

Notice that—just as Gale and Shapley (1962) showed—the matching produced in this way is stable, not just with respect to the strict preferences that follow step 0.1, but with respect to the underlying preferences elicited from the parties, which may have contained indifferences. That is, there can’t be a “blocking pair,” a student and a school, not matched to one another, who would prefer to be. The reason is that, if a student prefers some school to the one she was matched with in the algorithm, she must have already applied to that school and been rejected. This applies to the original preferences too, which may not be strict, since tie-breaking just introduces more blocking pairs; so any matching that is stable with respect to artificially strict preferences is also stable with respect to the original preferences. But those additional blocking pairs are constraints, and these additional constraints can harm welfare. A simple 1–1 (“marriage market”) matching example is sufficient to see what’s going on.

Example 1 (*Tie-breaking can be inefficient*). Let $M = \{m_1, m_2, m_3\}$ and $W = \{w_1, w_2, w_3\}$ be the sets of students and schools respectively, with preferences given by:

$$\begin{aligned} P(m_1) &= w_2, w_1, w_3 & P(w_1) &= [m_1, m_2, m_3] \\ P(m_2) &= w_1, w_2, w_3 & P(w_2) &= m_3, m_1, m_2 \\ P(m_3) &= w_1, w_2, w_3 & P(w_3) &= m_1, m_2, m_3 \end{aligned}$$

The brackets around w_1 ’s preferences indicate that w_1 is indifferent between any of $[m_1, m_2, m_3]$ while, in this example, everyone else has strict preferences. Since there is only one place at w_1 , but w_1 is the first choice of two students (m_2 and m_3), some tie-breaking rule must be used.

Suppose, at step 0 of the deferred acceptance algorithm, the ties in w_1 ’s preferences are broken so as to produce the (artificial) strict preference $P(w_1) = m_1, m_2, m_3$. The deferred acceptance algorithm operating on the artificial strict preferences produces $\mu_M = [(m_1, w_1); (m_2, w_3); (m_3, w_2)]$, at which m_1 and m_3 each receive their second choice (while m_2 receives his last choice). But note that the matching $\mu = [(m_1, w_2); (m_2, w_3); (m_3, w_1)]$ is Pareto superior for the students, as m_1 and m_3 each receive their first choice, so they are both strictly better off than at μ_M , and m_2 is not worse off. If the preferences of school

w_1 were in fact strict, the matching μ would be unstable, because m_2 and w_1 would be a blocking pair. But w_1 doesn't really prefer m_2 to m_3 ; in fact, μ is stable with respect to the original, non-strict preferences. The pair (w_1, m_2) is not a blocking pair for μ , and only appeared to be in the deferred acceptance algorithm because of the arbitrary ways in which ties were broken to make w_1 's preferences look strict.

So, there are costs to arbitrary or random tie-breaking. Erdil and Ergin (2006, 2007), Abdulkadiroğlu et al. (2009), and Kesten (2010) explore this from different angles.³⁹ Kesten notes that students are collectively better off at μ than at μ_M in example 1 because, in the deferred acceptance algorithm, m_2 's attempt to match with w_1 harms m_1 and m_3 without helping m_2 . Kesten defines an *efficiency-adjusted deferred acceptance mechanism* that produces μ in example 1 by disallowing the blocking pair (w_1, m_2) via a definition of "reasonable fairness" that generalizes stable matchings. But he shows that there is no mechanism that is Pareto efficient, reasonably fair, and strategy-proof.

To understand Erdil and Ergin's approach, note that the Pareto improvement from μ_M to μ in example 1 comes from an exchange of positions between m_1 and m_3 . This exchange doesn't introduce any new blocking pairs, since, among those who would like to change their positions, m_1 and m_3 are among the most preferred candidates of w_1 and w_2 . Since there weren't any blocking pairs to the initial matching, this exchange can occur without creating any new blocking pairs.

Formally, Erdil and Ergin define a *stable improvement cycle* starting from some stable matching to be a cycle of students who each prefer the school that the next student in the cycle is matched to, and each of whom is one of the school's most preferred candidates among the students who prefer that school to their current match. They prove the following theorem.

Theorem 15 (Erdil and Ergin, 2007). *If μ is a stable matching that is Pareto dominated (from the point of view of students) by another stable matching, then there is a stable improvement cycle starting from μ .*

This implies that there is a computationally efficient algorithm that produces stable matchings that are Pareto optimal with respect to students. The initial step of the algorithm is a student-proposing deferred acceptance algorithm with arbitrary tie-breaking of non-strict preferences by schools. The output of this process (i.e. the student optimal stable matching of the market with artificially strict preferences) is then improved by finding and satisfying stable improvement cycles, until no more remain. Erdil and Ergin show, however, that this algorithm is not strategy-proof; that is, unlike the student-proposing deferred acceptance algorithm, this deferred acceptance plus stable improvement cycle algorithm doesn't make it a dominant strategy for students to

³⁹ In the computer science literature there has been a focus on the *computational* costs of non-strict preferences, which adds to the computational complexity of some calculations (but not others) (see e.g. Irving, 1994; Irving et al., 2008). When preferences aren't strict, not all stable matchings will have the same number of matched people, and Manlove et al. (2002) show that the problem of finding a maximal stable matching is NP hard.

state their true preferences. They show in fact that no mechanism that always produces a stable matching that is Pareto optimal for the students can be strategy-proof.

Abdulkadiroğlu et al. (2009) establish that no mechanism (stable or not, and Pareto optimal or not) that is better for students than the student-proposing deferred acceptance algorithm with tie breaking can be strategy-proof. Following the design of the New York and Boston school choice mechanisms, define a *tie-breaking rule* T to be an ordering of students that is applied to any school's preferences to produce a strict order of students within each of the school's indifference classes (that is, when a school is indifferent between two students, the tie-breaking rule determines which is preferred in the school's artificial strict preferences). *Deferred acceptance with tie breaking rule* T is then simply the deferred acceptance algorithm operating on the strict preferences that result when T is applied to schools' preferences. One mechanism *dominates* another if, for every profile of preferences, the first mechanism produces a matching that is at least as good for every student as the matching produced by the second mechanism, and for some preference profiles the first mechanism produces a matching that is preferred by some students.

Theorem 16 (Abdulkadiroğlu et al., 2009). *For any tie-breaking rule T , there is no individually rational mechanism that is strategy-proof for every student and that dominates student-proposing deferred acceptance with tie-breaking rule T .*

But Abdulkadiroğlu, Pathak, and Roth also analyze the preferences submitted in recent NYC high-school matches (under a deferred acceptance with a tie-breaking mechanism) and find that, *if* the preferences elicited from the strategy-proof mechanism could have been elicited by a stable improvement cycle mechanism, then about 1,500 out of about 90,000 NYC students could have gotten a more preferred high school. (In contrast, the same exercise with the preferences submitted in the Boston school choice system yields almost no improvements.) So a number of open questions remain, among them, what accounts for the difference between NYC and Boston, and to what extent could the apparent welfare gains in NYC actually be captured? The potential problem is that, when popular schools are known, it's not so hard to find manipulations of stable improvement cycle mechanisms (which give families the incentive to rank popular schools more highly than in their true preferences, because of the possibility of using them as endowments from which to trade in the improvement cycles). Azevedo and Leshno (2010) show by example that at equilibrium such manipulations could sometimes be welfare decreasing compared to the (non-Pareto optimal) outcome of the deferred acceptance algorithm with tie-breaking.⁴⁰

So far I have been speaking of tie-breaking when a school is indifferent among a group of students only some of whom can be admitted. Students being indifferent among

⁴⁰ There has been a blossoming of new theory on school choice, including reconsideration of some of the virtues of the Boston algorithm, new hybrid mechanisms, and experiments. See for example Abdulkadiroğlu et al. (2010, 2011), Calsamiglia et al. (2010), Featherstone and Niederle (2010), Haeringer and Klijn (2009), Kojima and Unver (2010), and Mirrales (2009).

schools arose in a different way, because different seats in the same school (which are indistinguishable from the point of view of students) may be allocated according to different priority rules. We encountered this in New York because some schools, called Educational Option schools, are required to allocate half of their seats randomly, while the other half can be allocated according to the school's preferences. We also encountered it in Boston, where some schools use a "walk zone" priority for only half their seats. In each case, we created two "virtual schools" to which students could be admitted, one of which used each relevant priority rule. This is what introduced indifference in student preferences: each student was indifferent between a place in either of the virtual schools corresponding to a particular real school. But how these ties were broken could have consequences. So, for example, as reported in Abdulkadiroğlu et al. (2005a), the design decision we made in New York was that "If a student ranked an EdOpt school, this was treated in the algorithm as a preference for one of the random slots first, followed by a preference for one of the slots determined by the school's preferences." This was welfare improving for schools, since it meant that random slots would fill up before slots governed by the school's preferences, so a desirable student who happened to be admitted to a random slot would allow an additional preferred student to be admitted. However, other, more flexible rules can be considered. Kominers and Sönmez (2012) explore this issue with care, and reveal some subtle issues in the underlying theory.

New operational issues

One of the problems facing the old NYC school assignment system was congestion, caused in part by the time required for students who had received multiple offers to make a decision and allow waiting lists to move. In Boston, in contrast, the old school assignment system wasn't congested; it already used a centralized, computerized clearinghouse to give just one offer per student. Its problems arose from the way in which the assignment was made. However, as new kinds of public/private schools emerged, such as charter schools, Boston school choice has become something of a hybrid system, in which students get a single offer from the public school system but may get parallel offers from charter schools. Consequently, there is now some congestion and delay in processing waiting lists until these students choose which school to attend. Since the charter schools admit by lottery, this problem could easily be solved by including them in the centralized clearinghouse.

This is a problem we can hope to address from the outset as school choice technology continues to spread to other cities. Neil Dorosin, one of the NYC Department of Education administrators with whom we worked on the implementation of their high-school choice process, subsequently founded the non-profit Institute for Innovation in Public School Choice (IIPSC). With technical support from Abdulkadiroğlu, Pathak, and myself, IIPSC helped introduce new school choice systems in Denver and New Orleans. Denver uses a deferred acceptance algorithm, while in the Recovery School District in New Orleans the matching of children to schools in 2013 was due to be done by a version of a top trading cycles algorithm, along the lines discussed as a possibility for

Boston in Abdulkadiroğlu et al. (2005). The New Orleans school choice system includes charter schools (but not yet all of its schools).

Medical labor markets

Theoretical issues

One of the longstanding empirical mysteries regarding the medical labor market clearinghouse is why it works as well as it does in connection with helping couples find pairs of jobs. The story actually began sometime in the 1970s, when for the first time the percentage of women medical graduates from US medical schools rose above 10% (it is now around 50%). With this rise in women doctors came a growing number of graduating doctors who were married to each other, and who wished to find two residency positions in the same location. Many of these couples started to defect from the match. As noted in Roth (1984), not only does the deferred acceptance algorithm not produce a matching that is stable when couples are present (even when couples are allowed to state preferences over *pairs* of positions), but when couples are present it is possible that no stable matching exists. The following simple example from Klaus and Klijn (2005) makes this clear. This version is from Roth (2008b).

Example 2. *Market with one couple and no stable matchings (Klaus and Klijn 2005): Let $c=(s_1, s_2)$ be a couple, and suppose there is another (single) student, s_3 , and two hospitals, h_1 and h_2 . Suppose that the acceptable matches for each agent, in order of preference, are given by*

$$\begin{array}{ll} c: (h_1, h_2);^{41} & s_3: h_1, h_2 \\ h_1: s_1, s_3; & h_2: s_3, s_2 \end{array}$$

Then no individually rational matching μ (i.e. no μ that matches agents only to acceptable mates) is stable. We consider two cases, depending on whether the couple is matched or unmatched.

Case 1: $\mu(c) = (h_1, h_2)$. Then s_3 is unmatched, and he and h_2 can block μ , because h_2 prefers s_3 to $\mu(h_2)=s_2$.

Case 2: $\mu(c) = c$ (unmatched). If $\mu(s_3) = h_1$, then (c, h_1, h_2) blocks μ . If $\mu(s_3) = h_2$ or $\mu(s_3) = s_3$ (unmatched), then (s_3, h_1) blocks μ .

The new algorithm designed for the National Resident Matching Program by Roth and Peranson (1999) allows couples to state preferences over pairs of positions, and

⁴¹ Couple c submits a preference list over pairs of positions, and specifies that only a single pair, h_1 for student s_1 and h_2 for student s_2 , is acceptable. Otherwise couple c prefers to remain unmatched. For a couple, this could make perfect sense, if for example h_1 and h_2 are in a different city than the couple now resides, and they will move only if they find two good jobs.




FIGURE 1.2. High-level flowchart of the Roth and Peranson (1999) applicant-proposing deferred acceptance algorithm with couples.

seeks to find a stable matching (see Figure 1.2).⁴² The left side of the flow chart describes a fairly standard deferred acceptance algorithm with applicants proposing, much like the basic deferred acceptance algorithm described above in connection with school choice. However, because some applicants are couples who submit preferences over pairs of positions, it may be that a member of a couple sometimes needs to be *withdrawn* from a tentative assignment without having been displaced by a preferred applicant, something that never happens when all applicants are single. This occurs when one member of a couple is displaced by a preferred applicant, so the couple has to apply to another pair of positions, necessitating the withdrawal of the other couple member from the residency program that is holding his or her application. Since that residency program may have rejected other applicants in order to hold this one, this withdrawal may create blocking pairs. Therefore the right side of the flowchart describes an algorithm that tries to repair any blocking pairs that may have arisen in this way. Of course, the

⁴² The flowchart of the Roth-Peranson algorithm in Figure 1.2 was prepared for an early draft of Roth and Peranson (1999), but was removed in the editorial process, so it is published for the first time here (although it has been available on the Internet for some years in the lecture notes for my market design classes).

algorithm may cycle and fail to find a stable matching, as it must when there is no stable matching, for instance.

The empirical puzzle is why it almost never fails to find a stable matching, in the several dozen annual labor markets in which it has now been employed for over a decade (see Roth, 2008b, for a recent list). Some insight into this, reported in Kojima et al. (2010), connects the success in finding stable matchings that include couples to other recent results about the behavior of large markets.

Roth and Peranson (1999) initiated a line of investigation into large markets by showing computationally that if, as a market gets large, the number of places that a given applicant interviews (and hence the size of his rank order list) does not grow, then the set of stable matchings becomes small (when preferences are strict). Immorlica and Mahdian (2005) showed analytically that in a one-to-one marriage model with uncorrelated preferences, the set of people who are matched to different mates at different stable matchings grows small as the market grows large in this way, and that therefore the opportunities for profitable manipulation grow small. Kojima and Pathak (2009) substantially extend this result to the case of many-to-one matching, in which opportunities for employers to profitably manipulate can occur even when there is a unique stable matching, and in which employers can manipulate capacities as well as preferences. They show that as the size of a market grows towards infinity in an appropriate way, the proportion of employers who might profit from (any combination of) preference or capacity manipulation goes to zero in the worker-proposing deferred acceptance algorithm. Ashlagi et al. (2013) showed that small sets of stable matchings may be typical of large markets. Kojima et al. (2010) showed that when couples are present, if the market grows large in a sufficiently regular way that makes couples a small part of the market, then the probability that a stable matching exists converges to one. That is, in big enough markets with not too many couples we should not be surprised that the algorithm succeeds in finding a stable matching so regularly (see also Ashlagi et al., 2010).

A key element of the proofs is that if the market is large, but no applicant can apply to more than a small fraction of positions, then, even though there may be more applicants than positions, it is a high-probability event that there will be a large number of hospitals with vacant positions after the centralized clearinghouse has found a stable matching. This result is of interest independently from helping in the proofs of the results described above: it means that stable clearinghouses are likely to leave both people unmatched and positions unfilled, even when the market grows very large. Most clearinghouses presently have a secondary, post-match market, often called a “scramble,” at which these unmatched people and positions can find one another. The newly developing theory of large markets suggests that post-match marketplaces will continue to be important in markets in which stable centralized clearinghouses are used.

Operational issues

While there has been theoretical progress on managing post-match scrambles, some of this has yet to make its way into practice. In 2012 the National Resident Matching

Program introduced a formal scramble mechanism, called the Supplemental Offer and Acceptance Program. It appears to rely on punishments and sanctions to incentivize orderly participation, and my colleagues and I have expressed some reservations that this will be an effective design for the long term (Coles et al., 2010b).

The clearinghouse for gastroenterology fellowship positions discussed in the first part of this chapter seems to have established itself as a reliable marketplace; in the (2006) match for 2007 positions, 283 positions were offered and 585 applicants applied, of whom 276 were matched. In the match for 2011 positions, 383 positions were offered to 642 applicants, of whom 362 were matched (Proctor et al., 2011). This suggests that the policies adopted to decrease the frequency and effectiveness of exploding offers have been effective (see also Niederle and Roth, 2009a,b).⁴³ However Proctor et al. (2011) note that there are some warning signs that thickness may be difficult to maintain in the small part of the market that involves research positions. They observe that “the competition for these increasingly scarce, well-qualified, research-track applicants has become fierce, and the authors are aware of several examples during the last application cycle of candidates interested in research being offered fellowship positions outside the Match.”

Kidney transplantation

The theoretical and operational issues in kidney exchange are too intertwined for me to try to separate them here. Perhaps the most dramatic recent change in kidney exchange is that, following the publication of Rees et al.’s (2009a) report on the first non-simultaneous extended altruistic donor (NEAD) chain in the *New England Journal of Medicine*, there has been a small explosion of such chains, not only by established exchange networks, but also by transplant centers of all sorts around the United States. See for example the various chains reported at <<http://marketdesigner.blogspot.com/search/label/chains>>, or the more detailed report of chains conducted by the Alliance for Paired Donation (APD) in Rees et al. (2010). Simulations by Ashlagi et al. (2011a,b) using clinical data from the APD suggest that such chains can play an important role in increasing the number of live donor transplants, and recent theoretical progress has been made in understanding this in Ashlagi et al. (2012) (see also Ashlagi and Roth, 2012; and Dickerson et al., 2012).

The passage into law of what became the ‘Charlie W. Norwood Living Organ Donation Act’ (Public Law 110–144, 100th Congress) in December 2007 has set in motion plans that may eventually become a national kidney exchange network, but this is still moving slowly, and the issues involved with providing the right incentives for transplant centers to fully participate have not yet been resolved. Indeed, when I discussed this incentive problem in Roth (2008a) it looked like a problem that would become

⁴³ The job market for some other medical subspecialties continues to unravel, and orthopedic surgeons have recently taken steps to organize a centralized match (see Harner et al., 2008).

significant in the future, and today it has become a big issue. Ashlagi and Roth (2011) introduce a random graph model to explore some of these incentive issues in large markets, and show that the cost of making it safe for hospitals to participate fully is low, while the cost of failing to do so could be large if that causes hospitals to match their own internal patient–donor pairs when they can, rather than making them available for more efficient exchanges. That is, guaranteeing hospitals that patients whom they can transplant internally will receive transplants will not be too costly in terms of the overall number of transplants that can be accomplished in large markets. Among the easy-to-match pairs that hospitals withhold are those who are compatible, so that the donor can give directly to the intended recipient, even though such pairs might receive a better-matched kidney through exchange. The inclusion of compatible pairs would greatly increase the efficiency of kidney exchange, in no small part because it would ease the shortage of blood type O donors (see e.g. Roth et al., 2005; and Sönmez and Ünver, 2011; and see also Ünver, 2010, for a discussion of dynamic kidney exchange in large markets). But in the meantime, kidney exchange networks are seeing a disproportionate percentage of hard-to-match pairs, and Ashlagi et al. (2012) use models of sparse random graphs to suggest that this is costly in terms of lost transplants, and that it also accounts for why long ND donor chains have become so useful.

While kidney exchange is growing quickly⁴⁴ it still accounts for only a very small fraction of the number of transplants, and the growth is not yet enough to halt the growth of the waiting list for deceased-donor kidneys. (By early 2012 more than 90,000 candidates were registered on the kidney transplant waiting list in the United States.) This has led to continued discussion about ways to recruit more donors, and to continued interest in assessing views on whether kidneys might, in an appropriately regulated environment, under some circumstances be bought and sold, or whether donors could in some way be compensated. The whole question of compensation for donors remains an extremely sensitive subject.

For example, two recent surveys published in the surgical literature showed that public opinion and patient opinion both reflect a willingness to consider payment for organs (Leider and Roth, 2010; and Herold, 2010 respectively). However, the journal that published those surveys also published an editorial (Segev and Gentry, 2010) expressing the opinion that it was a waste of resources even considering the opinions of anyone other than physicians, and expressing the view that physicians were unalterably opposed to any change from current law prohibiting any “valuable consideration” for transplant organs. (This view of physician opinion seems not to be quite accurate, based on available surveys of physician opinion, and on the letters to the editor the journal

⁴⁴ See Wallis et al. (2011), with the caveat that the UNOS data on kidney exchange and ND donation appears to be incomplete, and may substantially underestimate the kidney exchange transplants to date, for instance because an initially ND donation may be recorded as a directed donation. The data collected by the US Department of Health and Human Services (Health Resources and Services Administration) at <<http://optn.transplant.hrsa.gov>> are incomplete and ambiguous, but suggest that between 367 and 636 transplants from exchange were reported to it in 2010, compared to between 228 and 441 in 2008, and between 34 and 121 in 2004. (The larger numbers come from including categories that today may include kidney exchange, but almost certainly did not in 2004.)

received in reply to what seems to be a fringe view.) Nevertheless, it is an indication that this remains a controversial subject, with views ranging widely, from those who might contemplate a fairly unregulated market (cf. Becker and Elias, 2007), to those who favor a moderately regulated market like the one in Iran (described in Fatemi, 2010), to those who would consider less direct forms of donor compensation (cf. Satel, 2009), to those, like the editorialists mentioned above, who consider the issue to be beyond discussion except insofar as it impacts physicians.

The continued shortage of kidneys (and other organs) for transplant therefore underlines the importance of continuing to try to expand deceased donation. Kessler and Roth (2012) report on possibilities of increasing donation by changing organ allocation policy to give increased priority to people who have been long-time registered donors. (This is an element of Singapore's organ allocation policy, and lately also Israel's policy.)

Economists and lawyers: two markets worth watching

Coles et al. (2010a) describe the recent experience of the market for new PhD economists with the newly instituted "pre-market" signaling mechanism, and "post-market" scramble. From 2006 through 2009, the number of candidates who used the signaling mechanism remained roughly constant at around 1,000 per year. The evidence is suggestive if not conclusive that judicious signaling increases the probability of receiving an interview. The pattern of signals suggests something about what might constitute "judicious" signaling; when one compares the reputational ranks of the school a student is graduating from and those he signals to, very few signals are sent from lower-ranking to higher-ranking schools. It appears that the signals play a coordination role in ameliorating congestion, with signals distributed across a very broad range of schools. Some new theory of "preference signaling" motivated by this market is presented in Coles et al. (forthcoming).

Participation in the post-market "scramble" has been more variable, with from 70 to 100 positions listed in each of the years 2006–10. It appears that at least 10% of these positions are filled each year through contacts made in the scramble.

Further developments in the market for new PhD economists will provide an ongoing window into the possibilities of dealing with congestion through signaling in a decentralized market, and in achieving thickness in the aftermarket.

A window of a different kind is being provided by several of the markets for new law graduates in the United States, which continue to suffer from problems related to the timing of transactions. The market for federal court clerks now appears to be nearing the end of the latest attempt to enforce a set of dates before which applications, interviews, and offers will not be made. Avery et al. (2007) already reported a high level of cheating in that market, as judges accepted applications, conducted interviews, and made offers before the designated dates. Roth and Xing (1994) reported on various ways that markets could fail through the unraveling of appointment dates, but the markets for lawyers have frequently offered the opportunity to observe new failures of

this kind. Presently the market for new associates at large law firms is also unraveling (see Roth, 2012).

Conclusions

The new marketplace designs reported in Roth (2008a), for labor markets, for schools, and for kidney exchange, have continued to operate effectively. However, in each of these domains, unsolved operational problems remain. In school choice, integrating standard public schools with other options such as charter schools in a single clearinghouse will help to avoid congestion. In kidney exchange, making it safe for hospitals to enroll all of their appropriate patient–donor pairs will help establish thickness and increase the number of transplants. In labor markets, it may be necessary to pay special attention to submarkets such as medical fellows interested in research.

These examples illustrate how market design, and the close attention it demands to the details of how particular markets operate, raises new theoretical questions about how markets work, and how market failures can be avoided and repaired. Holmstrom et al. (2002) quote Robert Wilson (1993) on this: “for the theorist, the problems encountered by practitioners provide a wealth of topics.”

REFERENCES

- 110th Congress, Public Law 110–144, Charlie W. Norwood Living Organ Donation Act, December 21, 2007, <http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=110_cong_public_laws&docid=f:publ144.110.pdf>.
- Abdulkadiroğlu, A. and Sönmez, T. (1999) “House allocation with existing tenants,” *Journal of Economic Theory*, 88: 233–60.
- (2003) “School choice: a mechanism design approach,” *American Economic Review*, 93(3): 729–47.
- Pathak, P. A. and Roth, A. E. (2005a) “The New York City high school match,” *American Economic Review, Papers and Proceedings*, 95(2): 364–7.
- Sönmez, T. (2005b) “The Boston public school match,” *American Economic Review Papers and Proceedings*, 95(2): 368–71.
- (2007) “Changing the Boston school choice mechanism: strategy-proofness as equal access,” NBER Working Paper No. 11965.
- (2009) “Strategy-proofness versus efficiency in matching with indifference: redesigning the NYC high school match,” *American Economic Review*, 99(5): 1954–78.
- Che, Y.-K., and Yasuda, Y. (2010) “Expanding ‘choice’ in school choice,” working paper.
- (2011) “Resolving conflicting preferences in school choice: the ‘Boston’ mechanism reconsidered,” *American Economic Review*, 101(1): 399–410.
- Abraham, D., Blum, A. and Sandholm, T. (2007) “Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges,” In *Proceedings of the ACM Conference on Electronic Commerce (EC)*.
- Ariely, D., Ockenfels, A. and Roth, A. E. (2005) “An experimental analysis of ending rules in internet auctions,” *Rand Journal of Economics*, 36(4): 891–908.

- Artemov, G. (2008) "Matching and price competition: would personalized prices help?" *International Journal of Game Theory*, 36(3): 321–31.
- Ashlagi, I. and Roth, A. E. (2011) "Individual rationality and participation in large scale, multi-hospital kidney exchange," Working Paper.
- (2012) "New challenges in multi-hospital kidney exchange," *American Economic Review: Papers and Proceedings*, 102(3): 354–9.
- Braverman, M. and Hassidim, A. (2010) "Matching with couples in large markets revisited," unpublished mimeo, MIT Sloan School.
- Gilchrist, D. S., Roth, A. E. and Rees, M. A. (2011a) "Nonsimultaneous chains and dominos in kidney paired donation—revisited," *American Journal of Transplantation*, 11(5): 984–94.
- (2011b) "NEAD chains in transplantation," *American Journal of Transplantation*, 11: 2780–1.
- Gamarnik, D., Rees, M. and Roth, A. E. (2012) "The need for (long) chains in kidney exchange," Working Paper.
- Kanoria, K. and Leshno, J. D. (2013) "Unbalanced Random Matching Markets," Working Paper.
- Avery, C., Jolls, C., Posner, R. A. and Roth, A. E. (2001) "The market for federal judicial law clerks," *University of Chicago Law Review*, 68: 793–902.
- Fairbanks, A. and Zeckhauser, R. (2003) *The Early Admissions Game: Joining the Elite*, Harvard University Press.
- Jolls, C., Posner, R. A. and Roth, A. E. (2007) "The new market for federal judicial law clerks," *University of Chicago Law Review*, 74: 447–86.
- Azevedo, E. M. and Leshno, J. D. (2010) "Can we make school choice more efficient? An incentives approach", draft, Harvard University.
- Becker, G. S. and Elías, J. J. (2007) "Introducing incentives in the market for live and cadaveric organ donations," *Journal of Economic Perspectives*, 21(3): 3–24.
- Biais, B., Hillion, P. and Spatt, C. (1999) "Price discovery and learning during the preopening period in the Paris Bourse," *Journal of Political Economy*, 107: 1218–48.
- Bulow, J. and Levin, J. (2006) "Matching and price competition," *American Economic Review*, 96(3): 652–68.
- Calsamiglia, C., Haeringer, G. and Klijn, F. (2010) "Constrained school choice: an experimental study," *American Economic Review*, 100(4): 1860–74.
- Cawley, J. (2006) "A guide (and advice) for economists on the U.S. junior academic job market," October <<http://www.aeaweb.org/joe/articles/2006/cawley.pdf>>.
- Chen, Y. and Sönmez, T. (2006) "School choice: an experimental study," *Journal of Economic Theory*, 127: 2002–231.
- Coles, P. A., Cawley, J. H., Levine, P. B., Niederle, M., Roth, A. E. and Siegfried, J. J. (2010a) "The job market for new economists: a market design perspective," *Journal of Economic Perspectives*, 24(4): 187–206.
- Featherstone, C. R., Hatfield, J. W., Kojima, F., Kominers, S. D., Niederle, M., Pathak, P. A. and Roth, A. E. (2010b) "Comment on the NRMP's "Supplemental Offer and Acceptance Program" proposed to replace the post-match scramble," <<http://kuznets.fas.harvard.edu/~aroth/papers/NRMP%20comment.pdf>>.
- Coles, P., Kushnir, A. and Niederle, M. (forthcoming) "Preference signaling in matching markets," *American Economic Journal: Microeconomics*.
- Cramton, P. (1997) "The FCC spectrum auctions: an early assessment," *Journal of Economics and Management Strategy*, 6(3): 431–95.

- Shoham, Y. and Steinberg, R. (eds) (2006) *Combinatorial Auctions*, MIT Press.
- Crawford, V. P. (2008) “The flexible-salary match: a proposal to increase the salary flexibility of the National Resident Matching Program,” *Journal of Economic Behavior and Organization*, 66: 149–60.
- Damiano, E., Li, J. and Suen, W. (2005) “Unraveling of dynamic sorting,” *Review of Economic Studies*, 72: 1057–76.
- Delmonico, F. L. (2004) “Exchanging kidneys: advances in living-donor transplantation,” *New England Journal of Medicine*, 350(18): 1812–14.
- Dickerson, J. P., Procaccia, A. D. and Sandholm, T. (2012) “Optimizing kidney exchange with transplant chains: theory and reality,” Working Paper.
- Elias, J. J. and Roth, A. E. (2007) “Econ one on one: kidney transplantation,” WSJ online, November 13 <http://online.wsj.com/public/article/SB118901049137818211.html?mod=todays_free_feature>.
- Erdil, A. and Ergin, H. (2007) “What’s the matter with tie-breaking? Improving efficiency in school choice,” *American Economic Review*, 98(3): 669–89.
- Ergin, H. and Sönmez, T. (2006) “Games of school choice under the Boston mechanism,” *Journal of Public Economics*, 90: 215–37.
- Evans, D. S. and Schmalensee, R. (with D.S. Evans) (1999) *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, MIT Press.
- Hagiu, A. and Schmalensee, R. (2006) *Invisible Engines: How Software Platforms Drive Innovation and Transform Industries*, MIT Press.
- Fatemi, F. (2010) “The regulated market for kidneys in Iran,” Sharif University of Technology, <http://gsme.sharif.edu/~ffatemi/Research/Kidney_Market-Farshad_Fatemi-11Jan2010.pdf>.
- Featherstone, C. and Niederle, M. (2010) “Ex ante efficiency in school choice mechanisms: an experimental investigation,” draft.
- Felli, L. and Sutton, J. (2006) “The Royal Economic Society’s first PhD presentations event,” <http://econ.lse.ac.uk/news/openfiles/js_RES_LSE_job_market_report.pdf>.
- Fréchette, G., Roth, A. E. and Ünver, M. U. (2007) “Unraveling yields inefficient matchings: evidence from post-season college football bowls,” *Rand Journal of Economics*, 38(4): 967–82.
- Gale, D. and Shapley, L. (1962) “College admissions and the stability of marriage,” *American Mathematical Monthly*, 69: 9–15.
- Gentry, S. E., Segev, D. L., Simmerling, M. and Montgomery, R. A. (2007) “Expanding kidney paired donation through participation by compatible pairs,” *American Journal of Transplantation*, 7: 2361–70.
- Grosskopf, B. and Roth, A. E. (2009) “If you are offered the right of first refusal, should you accept? An INVESTIGATION OF CONTRACT DESIGN,” *Games and Economic Behavior*, Special Issue in Honor of Martin Shubik, 65 (January): 176–204.
- Haeringer, G. and Klijn, F. (2009) “Constrained school choice,” *Journal of Economic Theory*, 144(5): 1921–47.
- Hanto, R. L., Roth, A. E., Ünver, M. U. and Delmonico, F. L. (2010) “New sources in living kidney donation,” in D. McKay (ed.), *Kidney Transplantation: A Guide to the Care of Transplant Recipients*, Springer, pp. 103–17.
- Saidman, S. L., Roth, A. E. and Delmonico, F. L. (2010) “The evolution of a successful kidney paired donation program,” XXIII International Congress of The Transplantation Society, August 16, Vancouver.

- Harner, C. D., Ranawat, A. S., Niederle, M., Roth, A. E., Stern, P. J., Hurwitz, S. R., Levine, W., DeRosa, G. P. and Hu, S. S. (2008) "Current state of fellowship hiring: Is a universal match necessary? Is it possible?" *Journal of Bone and Joint Surgery*, 90: 1375–84.
- Herold, D. K. (2010) "Patient willingness to pay for a kidney for transplantation," *American Journal of Transplantation*, 10: 1394–400.
- Holmstrom, B., Milgrom, P. and Roth, A. E. (2002) "Introduction to 'Game theory in the tradition of Bob Wilson,'" in B. Holmstrom, P. Milgrom and A. E. Roth (eds), *Game Theory in the Tradition of Bob Wilson*, Berkeley Electronic Press <<http://www.bepress.com/wilson>>.
- Immorlica, N. and Mahdian, M. (2005) "Marriage, honesty, and stability," *SODA*: 53–62.
- Irving, R. W. (1994) "Stable marriage and indifference," *Discrete Applied Mathematics*, 48: 261–72.
- Manlove, D. F. and Scott, S. (2008) "The stable marriage problem with master preference lists," *Discrete Applied Mathematics*, 156: 2959–27.
- Kessler, J. B. and Roth, A. E. (forthcoming) "Organ allocation policy and the decision to donate," *American Economic Review*.
- Jevons, W. S. (1876) *Money and the Mechanism of Exchange*, D. Appleton and Company.
- Kagel, J. H. and Levin, D. (2002) *Common Value Auctions and the Winner's Curse*, Princeton University Press.
- and Roth, A. E. (2000) "The dynamics of reorganization in matching markets: a laboratory experiment motivated by a natural experiment," *Quarterly Journal of Economics*, 115(1): 201–35.
- Kamecke, U. (1998) "Wage formation in a centralized matching market," *International Economic Review*, 39(1): 33–53.
- Kandel, E., Rindi, B. and Bosetti, L. (2007) "The effect of a closing call auction on market quality and trading strategies," Working Paper (Closing Call Auction in the Borsa Italiana).
- Kelso, A. S. and Crawford, V. P. (1982) "Job matching, coalition formation, and gross substitutes," *Econometrica*, 50(6): 1483–504.
- Kesten, O. (2010) "School choice with consent," *Quarterly Journal of Economics* 125(3): 1297–348.
- Klaus, B. and Klijn, F. (2005) "Stable matchings and preferences of couples," *Journal of Economic Theory*, 121(1): 75–106.
- Klemperer, P. (2004) *Auctions: Theory and Practice. The Toulouse Lectures in Economics*, Princeton University Press.
- Kojima, F. (2007) "Matching and price competition: comment," *American Economic Review*, 97(3): 1027–31.
- and Pathak, P. A. (2009) "Incentives and stability in large two-sided matching markets," *American Economic Review*, 99(3): 608–27.
- and Unver, M. U. (2010) "The 'Boston' school-choice mechanism," Working Paper, Boston College, February.
- Pathak, P. A. and Roth, A. E. (2012) "Matching with couples: stability and incentives in large markets," April 2010, revised September.
- Kominers, S. D. and Sönmez, T. (2012) "Designing for diversity in matching," Working Paper, September <http://www.scottkom.com/articles/Kominers_Sonmez_Designing_for_Diversity_in_Matching.pdf>.
- Lee, R. S. and Schwarz, M. (2007a) "Interviewing in two-sided matching markets," NBER Working Paper 14922.