

Using Expressiveness to Increase Efficiency in Social and Economic Mechanisms

Michael Benisch

May 2011
CMU-CS-11-113

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Norman Sadeh, co-chair
Tuomas Sandholm, co-chair
Geoff Gordon
Craig Boutilier (U. Toronto)

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy*

© 2011 Michael Benisch

This work has been supported by a Siebel Scholarship and National Science Foundation grants CNS-1012763, IIS-0964579, CNS-0905562, ITR-0905390, CNS-0627513, IIS-0427858, ITR-0205435, and IGERT-9972762. This research was also supported by CyLab at Carnegie Mellon under grants DAAD19-02-1-0389 and W911NF-09-1-0273 from the Army Research Office. Additional support has been provided by Nokia, France Telecom, Google, and the CMU/Portugal Information and Communication Technologies Institute.

Keywords: mechanism design, expressiveness, message space, efficiency, Vickrey-Clarke-Groves mechanism, combinatorial auction, multi-attribute auction, search keyword auction, sponsored search, catalog offers, menus, web commerce, usable privacy, location sharing, web services, social networking.

Abstract

Mechanisms for facilitating people’s interactions with businesses, their governments, and each other are ubiquitous in today’s society. One emerging trend over the past decade, along with increasing computational power and bandwidth, has been a demand for higher levels of expressiveness in such mechanisms. This trend has already manifested itself in combinatorial auctions and generalizations thereof. It is also reflected in the richness of preference expressions allowed by businesses as diverse as consumer sites, like Amazon and Netflix, and services like Google’s AdSense.

A driving force behind this trend is that greater expressiveness begets better matches, or greater efficiency of the outcomes. Yet, expressiveness does not come for free; it burdens users to specify more preference information. Today’s mechanisms have relied on empirical tweaking to determine how to deal with this and related tradeoffs. In this thesis, we establish the foundation of expressiveness in mechanisms and its relationship to their efficiency, as well as a methodology for determining the most effective forms of expressiveness for a particular setting.

In one stream of research, we develop a domain independent theory of expressiveness for mechanisms. We show that the efficiency of an optimally designed mechanism in equilibrium increases strictly as more expressiveness is allowed. We also show that in some cases a small increase in expressiveness can yield an arbitrarily large increase in a mechanism’s efficiency.

In a second stream of research, we operationalize our theory by applying it to a variety of domains. We first study a general class of mechanisms, called channel-based mechanisms, which subsume most combinatorial auctions. We show that without full expressiveness such mechanisms can be arbitrarily inefficient. Next, we focus on the domain of advertisement markets, where we show that the standard mechanism used for sponsored search is inefficient in the practical setting where some advertisers prefer lower-traffic positions (but this inefficiency can be largely eliminated by making the mechanism only slightly more expressive). We also consider the domain of privacy preferences for information sharing with one’s social network, where we conduct an extensive human subject study to determine which forms of expressiveness are most appropriate in the context of a location-sharing application. We conclude by developing and studying a framework for automatically suggesting high-profit prices in more expressive catalog pricing mechanisms (that allow sellers to offer discounts

on bundles in addition to pricing individual items). We use our framework to demonstrate several conditions under which offering discounts on bundles can benefit the seller, the buyer, and the economy as a whole.

To my soon-to-be wife and best friend Jessica Wisdom.

Acknowledgments

First and foremost, I would like to thank my advisors, Norman Sadeh and Tuomas Sandholm, whose guidance and wisdom were essential for finishing this thesis.

I am grateful to Norman for shaping me into the person I am today. It is because he continually pushed me to think about the practical implications of my work at an early stage of my career that I care so deeply about this aspect of research, a value that I believe to be critical in defining my professional and personal points of view. Norman’s door has always been open and I often took advantage of this fact to my great benefit. I would also like to thank him for exposing me to the important and interesting field of privacy and for spending numerous hours working with me on “wordsmithing” our papers until they were just right.

I am grateful to Tuomas for imparting his ability to keep the big picture in mind while maintaining a deep, meticulous attention to detail. The ability to excel on both of these levels simultaneously is a distinguishing feature of Tuomas’ work and I thank him for teaching me this skill. I would also like to thank Tuomas for the countless hours he spent with me writing and rewriting papers and thinking through the difficult aspects of proofs, algorithms, and methodologies. It is nearly impossible to overstate the amount he has contributed to my understanding of AI, computational thinking, and economics.

I am also thankful to both Norman and Tuomas for going above and beyond their responsibilities as advisors and providing immeasurable support and guidance in career-related and personal aspects of my life. Both have been full of encouragement and enthusiasm for my work, especially at times when others were less so. They hold themselves to incredibly high standards and motivate those around them to do the same. They have instilled in me a sense of confidence in myself and my writing. I am honored to have worked with them both.

I also especially want to thank the other members of my committee, Geoff Gordon and Craig Boutilier, for their incredibly helpful feedback and encouragement. From high-level discussions of the thesis in its early stages to low-level edits on the final text, both Geoff and Craig have been instrumental in the construction of this document.

Several other students and faculty at Carnegie Mellon have been especially helpful and influential. First, I cannot say enough about how much I learned from and relied heavily on the support of my close friends and colleagues at CMU, especially George Davis, Andrew Gilpin, Eric Daimler, Patrick Kelley, Abe Othman, and Sam Ganzfried. I also owe much thanks to the students who came before me, including Vince Conitzer, Kate Larson, Bradley

Malin, Edo Airoidi, and Ralph Gross. I am particularly thankful for many helpful discussions with other faculty at CMU, including Lorrie Cranor, Jason Hong, and Kathleen Carley, and for the opportunity to learn from many inspiring faculty, including Tom Mitchell, Carlos Guestrin, Guy Blelloch, Latanya Sweeney, Ziv Bar-Joseph, and Kathryn Roeder.

Special thanks belong to the professors and colleagues who inspired me to follow this path as an undergraduate. In particular, my thesis advisor at Brown, Amy Greenwald, was the first person to introduce to me to AI research and continues to be a mentor and friend. Others from Brown that I would like to thank include my friends, Dan Naftalis and Andy Hanflik, my colleagues in CS, Maggie Benthall, Katrina Ligett, Dan Licata, Michael Tschantz, and Victor Naroditskiy, and my teachers, Pascal Van Hentenryck, Michael Black, Philip Klein, John Hughes, and Shriram Krishnamurti.

I owe a great deal of thanks to a number of other students and researchers with whom I have had valuable discussions. While impossible to list them all here, I would like to thank Paul Hanks-Drielsma, Justin Cranshaw, David Pennock, William Walsh, Alberto Sardinha, Jialiu Lin, Ram Ravichandran, Eran Toch, Jana Diesner, Janice Tsai, Ian Fette, Jay Springfield, Justin Boyan, Rob Shields, David Thompson, Ariel Procaccia, David Pardoe, Chris Kiekintveld, Patrick Jordan, Troels Bjerre Sørensen, Alex Rogers, Ioannis Vetsikas, Wolf Ketter, Michael Wellman, David Parkes, Peter Stone, John Collins, Nicholas Jennings, Kevin Leyton-Brown, and Boi Faltings.

CMU's incredible administrators and support staff have also been important in helping me graduate. I'd like to thank Linda Francona, who helped coordinate almost every aspect of the logistics involved in my graduate studies over the last few years. Before Linda, Jennifer Lucas handled this task and I am extremely grateful to her as well. Other staff and administrators that have been essential in making this process smoother include Connie Herold, Marilyn Walgora, Catherine Copetas, Jim Skees, Ed Walter, and the SCS helpdesk.

I would also like to thank all of my New York and Pittsburgh friends for providing me with much needed social support. All of you know who you are, but I'd like to specifically thank Eric, Peter, Travis, Tony, Jesse, Holly, and Kate from New York, and Nora, Paul, Kristen, Lindsey, Carlos, Leith, Ben R., Drew, and Ben V. from Pittsburgh.

It would be impossible to ignore the support I received from my loving family. In particular, I would like to thank my dad and mom, David and Penny, my sister, Caryn, my stepmom, Bina, and my stepsiblings, Rae, Stephanie, and Doug. Special thanks to my dad for giving me my first Apple IIe computer and teaching me how to program it while I sat on his knee. I'd also like to thank my extended family, especially Bill, Maida, David, and Dara, and my soon-to-be in-laws, Starr, John, Jake, Becca, Jason, and Megan.

Last but definitely not least, I owe a huge debt of gratitude to my best friend and soon-to-be wife Jessica Wisdom. I cannot fathom where I would be today without your unending support and compassion. You fill my life with joy and happiness, and you have picked up more than your share of slack to make my journey through graduate school possible. Thank you; I dedicate this thesis to you.

Contents

1	Introduction	1
1.1	Thesis statement	4
1.2	Summary of contributions	4
1.3	Organization	6
2	A Theory of Expressiveness in Mechanisms	9
2.1	Introduction	10
2.2	Preliminaries	11
2.3	Characterizing the expressiveness of mechanisms	15
2.3.1	Impact-based expressiveness	16
2.3.2	Shattering-based expressiveness	18
2.3.3	Uses of the expressiveness measures	22
2.4	Expressiveness and efficiency	23
2.4.1	Conditions for full efficiency	25
2.4.2	The efficiency bound increases strictly with expressiveness	26
2.4.3	Inadequate expressiveness can lead to arbitrarily low efficiency for some preference distributions	26
2.4.4	Bayes-Nash implementation of the upper bound is always possible in private values settings	28

2.5	Expressiveness and communication complexity	32
2.6	Expressiveness in channel-based mechanisms	35
2.7	Conclusions and future research	41
2.8	Appendix: Proofs of all technical claims	45
3	Expressiveness in Advertisement Auctions	65
3.1	Introduction	66
3.2	Setting and background results	67
3.2.1	A summary of our expressiveness theory framework	68
3.3	Adapting our theory of expressiveness to ad auctions	69
3.4	The Premium GSP mechanism	71
3.5	Computing the efficiency bound	72
3.5.1	Integer programming formulation	72
3.5.2	Tree search for computing the bound	74
3.6	Experiments with GSP and PGSP	76
3.6.1	Experimental setup	77
3.6.2	Experiment 1: Varying agents' profit margin	79
3.6.3	Experiment 2: Varying agent diversity	81
3.7	Conclusions and future research	82
4	Expressiveness in Privacy Mechanisms	87
4.1	Introduction	88
4.2	Theoretical background	91
4.2.1	A general privacy mechanism model	91
4.2.2	Policy-based utility functions	93
4.2.3	Expressiveness and efficiency in privacy mechanisms	93
4.3	Methods for our user study	95

<i>CONTENTS</i>	xi
4.3.1 Study overview	96
4.3.2 Software	97
4.3.3 Privacy mechanisms we compare	99
4.3.4 Measuring accuracy with variable cost	103
4.3.5 Identifying privacy policies with user-burden considerations	105
4.4 Empirical findings from our user study	107
4.4.1 Survey results	108
4.4.2 Mobility patterns and preference statistics	110
4.4.3 Measuring the effects of different privacy mechanisms	112
4.4.4 Results related specifically to location sharing with advertisers	122
4.5 Conclusions and future research	126
4.6 Appendix: Survey materials	129
4.6.1 Pre-study survey	129
4.6.2 Post-study survey	132
5 Expressiveness in Price Catalogs	135
5.1 Introduction	136
5.2 Item pricing can be arbitrarily inefficient	138
5.3 Searching for profit maximizing prices	139
5.4 Estimating a rich customer valuation model	142
5.4.1 Deriving the maximum likelihood estimate	143
5.4.2 Fitting the valuation model to purchase data	144
5.5 Empirical results	146
5.5.1 Results with pricing algorithms	146
5.5.2 Results with the fitting algorithm	152
5.5.3 Results with a shopping cart generator	154
5.6 Conclusions and future research	156

6	Related work	159
6.1	Informational complexity	160
6.2	Work based on finding or characterizing equilibria	160
6.3	Expressiveness in dominant-strategy mechanisms	161
6.4	Applications of expressiveness in mechanisms	162
6.5	Related work on bundle pricing and CAs	163
6.6	Related work on location sharing	166
7	Conclusion	169
7.1	Review of contributions	171
7.2	Review of future work	174

List of Figures

2.1	By choosing between two expressions, $\theta_i^{(1)}$ and $\theta_i^{(2)}$, agent i can distinguish between the impact vectors $[A, B]$ and $[C, D]$ (enclosed in rectangles). The other agents are playing the pure strategy profile $[\theta_{-i}^{(x)}, \theta_{-i}^{(y)}]$	17
2.2	Channel-based representations of two auctions. The items auctioned are an apple (a) and an orange (o). The channels for each agent i are denoted x_i , y_i , and z_i . The possible allocations are A, B, C, and D. In each one, the items that Agent 1 gets are in the first braces, and the items Agent 2 gets are in the second braces.	37
2.3	An agent controlling non-overlapping sets of channels S_1 , S_2 , and S_3 can semi-shatter a pair of opponent profiles over A and B or C and D but not both. .	40
3.1	Part of the search tree for a distribution with 3 types, $[t_1^n, t_2^n, t_3^n]$, and 4 outcomes $[A, B, C, D]$. Circles represent internal nodes and squares represent leaf nodes. The dashed nodes are not expanded, but they would be considered by the algorithm. The expanded path corresponds to the assignment of $[A, C, C]$ to types t_1^n , t_2^n , and t_3^n , respectively.	74
3.2	Example of prototypical valuations for brand and value advertisers. The brand advertiser shown has $\mu = 1$ and the value advertiser has $\mu = 0.5$ (as defined in Table 3.1). Valuations are shown in expectation, not per click (these values will be negative when the amortized cost per click of running the site is high or the expected value of a conversion is low, which we vary in our experiments). Rank 0% means the bottom position and Rank 100% means the top position.	77

3.3	The value of the upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on four-agent instances. Results are averaged over 50 runs with different expected values for a conversion.	80
3.4	The value of the upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on three-agent instances. Results are averaged over 50 runs with different expected values for a conversion.	80
3.5	The value of our upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on four-agent instances. Large values of $E[\mu]$ correspond to runs in which higher ranking positions are more valuable for the value advertisers and vice versa.	82
3.6	The value of our upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on three-agent instances. Large values of $E[\mu]$ correspond to runs in which higher ranking positions are more valuable for the value advertisers and vice versa.	83
4.1	A screen shot of our web application displaying an example location on a map between 8:48pm and 9:02am.	100
4.2	A screen shot of an audit question asking whether or not a subject would have been comfortable sharing the location displayed on the map with the friends and family group. An audit question, like the one shown here, appeared below the map for each of the groups, at each location a subject visited. Drop down menus are only displayed because “Yes, during part of this time...” is selected.	101
4.3	Part of a search tree for identifying a subject’s most accurate privacy policy using the Loc/Time+ mechanism.	106
4.4	The average percentage of time shared with each group during each thirty-minute interval throughout the day on weekdays (top) and weekends (bottom).	111
4.5	The average accuracy (bars indicate 95% confidence intervals) for each group, under each of the different privacy mechanisms. For these results, we hold constant the cost for inappropriately revealing a location at $c = 20$	114

4.6	The average accuracy for the Facebook friends group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.	116
4.7	The average percentage of time shared (bars indicate 95% confidence intervals) with each group under each of the different privacy mechanisms. For these results, we hold constant the cost for inappropriately revealing a location at $c = 20$	117
4.8	The average percentage of time shared with the Facebook friends group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.	118
4.9	The average accuracy (vertical axis) achieved by each of the different privacy mechanisms, for each of the different groups, varying a global limit on the number of rules (horizontal axis, from one to five or more) in a policy. We hold constant the cost for inappropriately revealing a location at $c = 20$, and identify policies with the highest possible total accuracy across all groups, while weighting each group equally.	120
4.10	The average accuracy (bars indicate 95% confidence intervals) achieved by each of the different privacy mechanisms for the Facebook friends group, while varying a limit on the number of rules in a policy that apply to Facebook friends only. We hold constant the cost for inappropriately revealing a location at $c = 20$	121
4.11	User responses on qualities of advertisers which would impact their future location-sharing decisions. Answers were reported on a 7-point Likert scale, from 1 (Not important) to 7 (Very, very important). Averages and 95% confidence intervals are shown.	123
4.12	The average percentage of time shared with the Advertiser group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.	124

- 4.13 The average accuracy (bars indicate 95% confidence intervals) achieved by each of the different privacy mechanisms for the Advertisers group, while varying a limit on the number of rules in a policy that apply to advertisers only. Results are shown for two different values of c , $c = 20$ and $c = 100$ 125
- 5.1 Illustration of the search technique we use for estimating a customer valuation model from historical purchase data. We use a tree search over variances and a pivot-based search over means. Leaves are evaluated based on how closely the corresponding model predicts the observed data and (optionally) how closely the model's optimal profit matches the profit achieved by existing prices. 147
- 5.2 The intensity of each dot is the increase in expected profit or surplus achieved by profit-maximizing bundle discounts for different levels of covariance (x-axis) and complementarity (or substitutability) (y-axis), ranging from 0% to 10%. Here, we assume the seller holds the item prices fixed at the optimal item-only catalog values to isolate the impact of bundle discounts. 150
- 5.3 The intensity of each dot is the increase in expected profit or surplus achieved by profit-maximizing bundle discounts for different levels of covariance (x-axis) and complementarity (or substitutability) (y-axis). Here, we assume the seller has the ability to reprice items as well as offer discounts on bundles. White dots on the surplus graph in this figure indicate that the surplus after repricing was the same or worse than the surplus under the optimal item-only catalog (this is done for consistency with Figure 5.2). 151
- 5.4 The intensity of each dot represents the predicted increase in expected profit or surplus achieved by profit-maximizing bundle discounts on single-catalog instances with varying item frequencies (x-axis) and co-occurrence percentages (y-axis), ranging from 0% to 10%. As in Figure 5.2, we assume the seller holds the item prices fixed at the optimal item-only catalog values to isolate the impact of offering bundle discounts from the potential confound of our system improving the item pricing as well. 153

List of Tables

3.1	Default settings for each parameter in our instance generation model.	79
4.1	An illustration of the policy-based utility function class under each of the four possible scenarios: i) the mechanism correctly grants, ii) correctly denies, iii) inappropriately grants or iv) inappropriately denies.	94
4.2	The average report on our pre-study survey of how comfortable subjects would have been on a 7-point Likert scale from “not comfortable at all” to “fully comfortable” if their location could be checked by each of the groups “Any-time,” “At locations you have specified,” or “At times you have specified.” .	109
4.3	The average report of how bad subjects thought it would have been, on a 7-point Likert scale from “not bad at all” to “very, very bad,” if their location were mistakenly withheld from or revealed to each of the groups.	109
4.4	The average percentage of time a subject spent at his or her five most visited locations, and the average percentage of time he or she would have shared that location with friends and family (FF), Facebook friends (FB), university community (UC), and advertisers (AD).	112
5.1	An illustration of all of the gradients considered by the pivot algorithm during each step with two products and $k = 2$. Each group of three cells in the table (enclosed in angle brackets) represents a gradient, an arrow indicates the direction of the corresponding bundle’s price movement. A cell with no arrow indicates no movement in the corresponding bundle’s price.	142

5.2	The average fraction of the highest expected profit, efficiency, and surplus, as well as the average number of calls to \hat{P} for each of the pricing algorithms described in Section 5.3. For these results, the algorithms were run on symmetric two-item instances.	148
5.3	The total profit and surplus increases for various parameter settings of a classic market basket analysis generator [5] (values are averaged over five instances for each parameter setting). Here, we assume the seller had priced all items optimally, and at those prices each item was sold for a uniform profit. We also consider discounts only on pairs of items that are unrelated to any other items.	155

Chapter 1

Introduction

Mechanism design is the science of generating rules of interaction so that desirable outcomes result despite the participating agents (human or computational) acting based on rational self-interest. A *mechanism* takes as input some expressions of preference from the agents, and based on that information imposes an *outcome* (such as an allocation of items and potentially also payments). By carefully crafting mechanisms, it is possible to design better auctions, exchanges, catalog offers, voting systems, privacy enforcing mechanisms, and so on.

Mechanisms that facilitate the interactions people have with businesses, their governments, and each other are ubiquitous in today’s society. One emerging trend over the past decade, along with increasing computational power and bandwidth, has been a demand for higher levels of expressiveness in mechanisms that mediate interactions such as the allocation of resources, matching of peers, and elicitation of opinions. This trend has already manifested itself in combinatorial auctions and generalizations thereof. It is also reflected in the richness of preference expression offered by businesses as diverse as consumer sites with product ratings, like Amazon and Netflix, and services like Google’s AdSense. In Web 2.0 parlance, the demand for increasingly diverse offerings is often referred to as the Long Tail [6].

The most famous expressive mechanism is a *combinatorial auction (CA)*, which allows participants to express valuations over *packages* of items in addition to valuations over the items themselves. CAs have the recognized benefit of removing the “exposure” problems that bidders face when they have preferences over packages but in traditional auctions are allowed to submit bids on individual items only. They also have other acknowledged benefits, and preference expression forms significantly more compact and more natural than package bidding have been developed (e.g., [32, 55, 73, 110, 125, 130, 132]). Expressiveness also plays a key role in *multi-attribute* settings where the participants can express preferences over vectors of attributes of the item—or, more generally, of the outcome. Some market designs are both combinatorial and multi-attribute (e.g., [55, 125, 130, 132]). Other examples of mechanisms that have become more expressive recently include e-commerce sites that have expanded their catalog offerings with bundles of items sold together (often accompanied by discounts), online advertisement auctions that allow advertisers to target their ads to particular geographical locations, and in-depth privacy control mechanisms for popular social networking web sites such as Facebook and LinkedIn.

Intuitively, allowing people or organizations to express richer preferences should yield more efficient outcomes (i.e., with higher average utilities of the participants). For example, enabling businesses to specify more expressive preferences in sourcing auctions has been shown to produce outcomes of significantly higher efficiency, leading to savings of billions of dollars each year (e.g., [55, 123, 125, 130, 131]). However, increasing expressiveness does not always improve matters. In some cases, increased expressiveness can backfire, leading to reduced competition and revenue, or confusion among the mechanism’s human participants who may not always act fully rationally [135, 136]. Furthermore, expressing complex preferences (e.g., asking businesses to evaluate a large number of possible supplier arrangements or asking a user to express preferences across a wide range of computer configurations) can be resource-intensive, costing companies money or users time [54, 122, 134]. A large body of research exists to address the issue of how to elicit such preferences in ways that minimize the number of user queries required, which is referred to as the *preference elicitation problem* (e.g., [26, 30, 31, 34, 88, 111, 126]). This work is complementary to ours because it aims to unlock the benefits of more expressive mechanisms without any unnecessary additional user burden.

Until now, we have lacked a general way of characterizing the expressiveness of different mechanisms, the impact that it has on the agents’ strategies, and thereby ultimately the outcome. For example, prior to our work, it was not even known whether, in any domain, more expressiveness could always be used to design economic mechanisms with more efficient equilibria. (In fact, in certain settings it had been shown that additional expressiveness can give rise to additional equilibria of poor efficiency [98].) Short of empirical tweaking, participants in the scenarios we described lack results they can rely on to determine how much—and what forms of—expressiveness they need. These questions have vexed mechanism design theorists, but are not only theoretical in nature. Answers could ensure that ballots are expressed in a form that matches the issues voters care about, that companies are able to identify suppliers that best match their needs, that supply and demand are better matched in B2C and C2C markets, that users of online social-networking sites can express those privacy preferences that really matter, and so on.

1.1 Thesis statement

It is possible to improve the efficiency of a wide variety of social and economic mechanisms, in theory and in practice, by using a computational framework for designing them with the most appropriate empirically determined levels and forms of expressiveness.

1.2 Summary of contributions

In Chapter 2, we begin by developing a **theoretical framework** [20, 23] that characterizes the impact of a mechanism’s expressiveness on its outcome in a domain-independent manner. As part of this work, we introduce two new notions of expressiveness, impact dimension and outcome shattering, based on ideas from computational learning theory. Our main results prove that a mechanism designer can strictly increase expected efficiency by giving any agent more expressiveness (until reaching full efficiency). Furthermore, we prove that this can be accomplished with a budget-balanced, Bayes-Nash incentive compatible mechanism (where participants are incentivized to reveal their true valuations in expectation), but we also show that, without full expressiveness, it cannot always be accomplished with a mechanism that is dominant-strategy incentive compatible (where participants are incentivized to reveal their true preferences no matter what). We then apply this general framework to a specific class of mechanisms, which we call channel based, and show that any (channel-based) multi-item auction without rich combinatorial bids can be arbitrarily inefficient.

In the remainder of the dissertation, we operationalize our theoretical framework by developing a **methodology** to compare mechanisms with different degrees and forms of expressiveness in different application domains. At a high level, the methodology, which uses a variety of models, algorithms, and techniques, involves i) estimating preference distributions for participants in a target domain, ii) identifying mechanisms that represent different degrees and forms of expressiveness, iii) computing socially optimal, equilibrium, or heuristic strategies for the agents under each of the mechanisms, iv) simulating the outcomes under the strategies that were computed, and v) comparing the outcomes based on, for example, their expected efficiency.

The first application area we explore (Chapter 3) is that of **advertisement markets**

[21]. These markets account for over \$200 billion in annual revenue across all media, and involve some of the fastest-growing mechanisms on the Internet. The most popular online ad mechanism, the generalized second price (GSP) mechanism used by Google, Yahoo!, Bing, Baidu, and others, solicits a single bid from each advertiser for a particular keyword, and assigns advertisers to positions on search-result pages according to these bids. We prove that, since it does not allow advertisers to express different bids for different positions, the GSP is inexpressive according to our domain-independent notions of expressiveness and, consequently, can be arbitrarily inefficient for some preference distributions. However, we also propose a new mechanism, called the Premium GSP (PGSP), which involves a small, intuitive increase in expressiveness by soliciting a single extra bid from each advertiser (the extra bid is for the right to appear in a premium position). Our empirical results, which involve simulating cooperative and heuristic strategies for the bidders, demonstrate that the PGSP can remove the bulk of the GSP's inefficiency in many realistic settings, which can be up to 30%. Concurrent with our work, Google adopted a feature similar to our premium mechanism, called position preference, suggesting that this type of mechanism is also useful in practice.

The second application area we consider (Chapter 4) is **privacy** [19, 85, 116]. The past few years have seen an explosion in the range of websites allowing individuals to exchange personal information and content that they have created. These sites include location-sharing services, social-networking services, and photo- and video-sharing services. While there is clearly a demand for people to share this information with each other, there is also a substantial demand for greater expressiveness in the privacy mechanisms that control how the information is shared. To apply our methodology in this domain, we performed a three-week user study in which we tracked the locations of 27 subjects and asked them to rate when, where, and with whom they would have been comfortable sharing their locations. Using the detailed preferences we collected, we identify the best possible policy (or collection of rules granting access to one's location) for each subject and privacy mechanism. To quantify the effects of different levels and forms of expressiveness, we measure the accuracy with which the resulting policies are able to capture our subjects' preferences. We also vary our assumptions about the sensitivity of the information and users' tolerance for the added burden associated with making more complex policies. Our results reveal that many of today's location-sharing applications, such as Loopt and Google's Latitude, may have failed to gain traction due to

their limited privacy settings.

In Chapter 5, we investigate a third and final application area of **catalog pricing** [22]. Business to customer retail sales account for nearly four trillion dollars in the United States annually, and the percentage of this shopping done online increased more than three-fold from 2002 to 2007. Yet, despite the increased computational power, connectivity, and data available today, most online and brick-and-mortar retail mechanisms remain nearly identical to their centuries-old original form (i.e., catalog pricing with take-it-or-leave-it offers). This is the default mechanism for brick-and-mortar B2C trade and is used by massive online retailers like Amazon, BestBuy, and Dell. In the final chapter of this dissertation, we begin to develop advances toward more expressive catalog pricing mechanisms that could thus lead to significant efficiency improvements across the economy. First, we show that our theoretical framework for studying expressiveness can be used to characterize the inefficiency of a commonly used inexpressive mechanism: the item-only catalog (i.e., a traditional catalog that offers prices for individual items only). We then describe a set of general algorithms for identifying profit-maximizing prices that repeatedly query a customer demand distribution with different candidate catalogs. We provide a method for learning this demand distribution from data, a task that we show is similar to the classic market basket analysis problem. (Market basket analysis involves counting the frequencies of different item sets and has been extensively studied, including by Google co-founders Sergey Brin and Larry Page [35, 36]). Finally, we perform computational experiments using our pricing and fitting algorithms to demonstrate several conditions under which offering discounts on bundles can benefit the seller, the buyer, and the economy as a whole.

1.3 Organization

The rest of this thesis is organized into four main chapters, each of which covers one of the broad topics outlined in the summary of contributions above (other than the general methodology, which is touched on in all of the chapters). Each chapter includes i) an introduction to the topic, ii) a description of all the methods developed on the topic for this dissertation, iii) a discussion of the results of applying those methods to simulated and (in some cases) real-world data, and iv) a conclusion and discussion of future work on the topic. We discuss related work throughout the dissertation and describe some of the most

closely related work in more detail in Chapter 6. Chapter 7 summarizes the dissertation and provides some concluding remarks.

Chapter 2

A Theory of Expressiveness in Mechanisms

2.1 Introduction

In this chapter, we develop a theory that ties the expressiveness of mechanisms to their efficiency in a domain-independent manner. We begin in Section 2.3 by introducing two notions of expressiveness: i) *impact dimension*, which captures the extent to which an individual agent can impact the mechanism’s outcome, and ii) *outcome shattering*, which is based on the concept of *shattering*, a measure of functional complexity from computational learning theory. We refer to increases (or decreases) in these measures as increases (or decreases) in expressiveness.

In Section 2.4, we derive an upper bound on the expected efficiency of any mechanism under its most efficient Bayes-Nash equilibrium. (In a Bayes-Nash equilibrium no agent can gain in expectation by unilaterally deviating.) This allows us to sidestep two of the major roadblocks in analyzing the relationship between expressiveness and efficiency: 1) the bound can be studied without having to solve for any of the mechanism’s equilibria (which tends to be extremely difficult for inexpressive mechanisms, e.g., [103, 107, 119, 139, 155, 159]), 2) since it bounds the *most efficient* equilibrium it can be used to study mechanisms with multiple—or even an infinite number of—equilibria, e.g., first price CAs [24]. Additionally, as we will show, a mechanism can incentivize agents to play the strategies prescribed by this bound in Bayes-Nash equilibrium by acting like a moderator.

We show that in any setting the bound of an optimally designed mechanism increases *strictly* as more expressiveness is allowed and, for some distributions over agent valuations, by an arbitrarily large amount via a small increase in expressiveness. We also prove that in any private values setting (i.e., where an agent’s utility depends only on its own private information and not the private information of any other agent) the bound is tight in that it is always possible to achieve its efficiency with a budget balanced mechanism in Bayes-Nash equilibrium. Taken together, these results imply that for any private values setting the expected efficiency of the best Bayes-Nash equilibrium increases strictly as more expressiveness is allowed. Interestingly, unlike with full expressiveness, implementing this bound is not always possible in dominant strategies. (In a dominant-strategy equilibrium no agent can gain by deviating, no matter what the other agents do.) Additionally, the efficiency of the bound may not be achieved by a mechanism if its payment function is not properly designed to incentivize it. Still, these results provide a significant step forward in our understanding

of the relationship between expressiveness and efficiency.

In Section 2.5, we explore the relationship between our expressiveness measures and more traditional notions of communication complexity, such as the amount of information agents must transmit. Specifically, we show that our expressiveness measures can be used to derive both upper and lower bounds on the number of bits needed by the best multi-party communication protocol for computing a given outcome function.

Finally, we study a class of mechanisms that we call *channel based*. They subsume most combinatorial allocation mechanisms (of which CAs and multi-attribute auctions are a subset) and any *Vickrey-Clarke-Groves (VCG)* scheme [44, 65, 151]. We show that our domain-independent measures of expressiveness appropriately relate to the natural measure of expressiveness for channel-based mechanisms: the number of channels allowed (which itself generalizes a classic measure of expressiveness in CAs called *k-wise dependence* [51]). Using this bridge, our general results yield interesting implications. For example, we prove that for any (channel-based) combinatorial allocation mechanism that does not allow rich combinatorial bids there exist distributions over agent valuations (even distributions satisfying the *free disposal* condition, i.e., where the utility of winning an extra item is always non-negative), for which the mechanism cannot achieve 95% of optimal efficiency. This 5% inefficiency is an order of magnitude greater than a related inefficiency previously proven for combinatorial allocation mechanisms with sub-exponential communication [105].

2.2 Preliminaries

The setting we study in this chapter is that of standard mechanism design. In the model there are n agents. Each agent, i , has some private information (not known by the mechanism or any other agent) denoted by a type, t_i (e.g., the value of the item to the agent in an auction; or, in a CA, a vector of values, potentially one for each bundle). The space of an agent's possible types is denoted T_i . We use the notation t^n to refer to a collection of n types (we occasionally omit the n superscript when it is clear that the entity is a collection of n types). Agent i 's types are drawn according to some distribution, $P(t_i)$, that we assume is known to the mechanism designer and to agent i , but not necessarily to all agents.

Each agent has a valuation function, $v_i(o, t_i)$, that indicates its valuation under type t_i ,

or how much utility the agent gets when it draws type t_i and outcome $o \in \mathcal{O}$ is chosen. We call the distribution of utilities defined by a valuation function and a corresponding probability distribution over types a *preference distribution*. Settings where each agent's valuation function depends only on its own type and the outcome chosen by the mechanism (e.g., the allocation of items to the agent in a CA) are called *private values* settings. We also discuss more general *interdependent values* settings, where $v_i = v_i(o, t^n)$ (i.e., an agent's valuation depends on the others' private signals). In both types of settings, agents report expressions to the mechanism, denoted θ_i , based only on their own types. We use the notation θ^n to refer to a collection of n expressions. A mapping from types to expressions is called a *pure strategy*.

Definition 1 (pure strategy). *A pure strategy for an agent i is a mapping, $h_i : T_i \rightarrow \Theta_i$, that selects an expression for each of i 's types. A pure strategy profile for a subset of agents, I , is a list of pure strategies, one strategy per agent in I , i.e., $h_I \equiv [h_1, h_2, \dots, h_{|I|}]$. For shorthand, we often refer to h_I as a mapping from types of the agents in I to an expression for each agent, $h_I(t_I) = [\theta_1, \theta_2, \dots, \theta_{|I|}]$.*

We also consider *mixed strategies*, or mappings from types to random variables specifying probability distributions over possible expressions.

Definition 2 (mixed strategy). *A mixed strategy for agent i is a mapping, $h_i : T_i \rightarrow P(\Theta_i)$, that selects a probability distribution over expressions for each of i 's types. A mixed strategy profile is a list of mixed strategies, one strategy per agent.*

Based on the expressions made by the agents, the mechanism computes the value of an outcome function, $f(\theta^n)$, which chooses an outcome from \mathcal{O} . The mechanism may also compute the value of a payment function, $\pi_i(\theta^n)$, which determines how much each agent, i , must pay or get paid.¹

In Section 2.4, we discuss results pertaining to the implementation of a mechanism under two different solution concepts: Bayes-Nash and dominant strategy equilibria. We do not re-

¹In Section 2.3, we define our measures of expressiveness based only on the mechanism's outcome function. For our purposes, this is without loss of generality as long as agents do not care about *each others'* payments. We later discuss the payment function in more depth when we examine issues related to incentives in Section 2.4.

strict our attention to mechanisms with truthful equilibria (i.e., where agents are incentivised to report their true types in equilibrium).²

Definition 3 (Bayes-Nash equilibrium). *A strategy profile is a Bayes-Nash equilibrium when no agent can gain expected utility by unilaterally deviating (i.e., assuming the expressions of all the other agents remain fixed). Formally, a (potentially) mixed strategy profile, m , constitutes a Bayes-Nash equilibrium for outcome function f and payment function π if*

$$\forall i, \forall m'_i \neq m_i, \int_{t^n} P(t^n) \int_{\theta^n} P(m(t^n) = \theta^n) u_i(f(\theta^n), t_i) - \pi_i(\theta^n) \geq \int_{t^n} P(t^n) \int_{\theta^n} P(\{m'_i(t_i), m_{-i}(t_{-i})\} = \theta^n) u_i(f(\theta^n), t_i) - \pi_i(\theta^n)$$

Definition 4 (dominant-strategy equilibrium). *A pure-strategy profile is a dominant-strategy equilibrium when no agent can gain utility by deviating, regardless of how many other agents also do so. Formally, a pure-strategy profile, h , constitutes a dominant-strategy equilibrium for outcome function f and payment function π if*

$$\forall i, \forall h'_i \neq h_i, \forall t^n, u_i(f(h(t^n)), t_i) - \pi_i(h(t^n)) \geq u_i(f(h'_i(t_i), h_{-i}(t_{-i})), t_i) - \pi_i(h'_i(t_i), h_{-i}(t_{-i}))$$

During some of our analysis, we consider the widely studied class of mechanisms in which the set of expressions available to an agent corresponds directly with its types. These are called *direct-revelation mechanisms*.

Definition 5 (direct-revelation mechanism). *A direct-revelation mechanism is a mechanism in which each agent's expression space is equal to its type space (i.e., $T_i = \Theta_i$, for all i).*

To summarize, we use the following notation.

- $t_i \in T_i$ is the true type of an agent i . The subscript t_{-i} is used to denote a set of types for all the agents other than i , and the superscript t^n is used to denote a set of n types.

²The *revelation principle* of mechanism design states that any outcome function that can be implemented by any mechanism under a non-truthful equilibrium can also be implemented by some mechanism under a truthful equilibrium [94]. However, we do not restrict our analysis to mechanisms with truthful equilibria because in mechanisms without full expressiveness it can be impossible for agents to express their true types.

- $\theta_i \in \Theta_i$ is the expression that agent i reports to the mechanism. The subscript θ_{-i} is used to denote a set of expressions for all the agents other than i , and the superscript θ^n is used to denote a set of n expressions.
- $o \in \mathcal{O}$ is an outcome from the set of all possible outcomes imposable by the mechanism.
- $v_i : \mathcal{O}, T_i \rightarrow \mathbb{R}$ is agent i 's valuation function. It takes as input the agent's true type and an outcome and returns the real-valued utility of the agent if that outcome were to be chosen. (We also discuss results that apply to interdependent values settings, where $v_i = v_i(o, t^n)$, i.e., an agent's utility also depends on others' private signals.)
- $f : \Theta^n \rightarrow \mathcal{O}$ is the outcome function of the mechanism. It takes as input the expression of each agent and returns an outcome from the set of all possible outcomes.
- $\pi : \Theta^n \rightarrow \mathbb{R}^n$ is the payment function of the mechanism. It takes as input the expression of each agent and returns the payment to be made by each agent.

For convenience, we will let $W(o, t^n)$ denote the total social welfare of outcome o when agents have private types (or private signals) t^n , i.e., $W(o, t^n) = \sum_i v_i(o, t^n)$. Occasionally, we use the shorthand W_I , where I refers to some subset of the agents, to denote the total social welfare of only the agents in I . Assuming the agents play a mixed strategy profile denoted by m , the expected efficiency, $E[\mathcal{E}(f)]$, of an outcome function, f , (where expectation is taken over the types of the agents and their randomized equilibrium expressions) is given by

$$(2.1) \quad E[\mathcal{E}(f)] = \int_{t^n} P(t^n) \int_{\theta^n} P(m(t^n) = \theta^n) W(f(\theta^n), t^n).$$

The following example shows how this formalism can be used to model a combinatorial auction.

Example 1. *In a fully expressive combinatorial auction with m items, each of the agents is a bidder whose type represents his or her private valuation for each of the 2^m different combinations of items. The outcome space includes all of the n^m different ways the items can be allocated among the bidders. Agents are allowed to express their entire type to the mechanism and the outcome function chooses the allocation that maximizes the sum of the bidders' valuations.*

The payment function can charge each agent its bid (a.k.a. the first-price payment rule) or the difference in utility of the other agents had the agent in question not participated (a.k.a. the Vickrey-Clarke-Groves (VCG) payment rule). Under the VCG payment rule, each agent has a (weakly) dominant strategy to tell the truth, so one equilibrium distribution over expressions is a point mass on the agents' true valuations.

2.3 Characterizing the expressiveness of mechanisms

The primary goal of this chapter is to better understand the impact of making mechanisms more or less expressive. In order to achieve this goal, we must first develop meaningful (and general) measures of a mechanism's expressiveness. We will begin by demonstrating that two seemingly natural ways of characterizing the expressiveness of different mechanisms, the dimensionality of their expressions and the granularity of their outcomes, do not capture the fundamental difference between expressive and inexpressive mechanisms. Later, in Section 2.5, we will discuss the relationship between expressiveness and communication complexity, which can be thought of as the granularity of the expression space.

If we consider mechanisms that allow expressions from the set of multi-dimensional real numbers, such as CAs and combinatorial exchanges, one seemingly natural way of characterizing their expressiveness is the dimensionality of the expressions they allow (e.g., this is one difference between CAs and auctions that only allow per-item bids). However, not only would this limit the notion of expressiveness to mechanisms with real-valued expressions, it also does not adequately differentiate between expressive and inexpressive mechanisms, as the following well-known result demonstrates.

Proposition 1. *For any mechanism that allows multi-dimensional real-valued expressions, (i.e., $\Theta_i \subseteq \mathbb{R}^d$), there exists an equivalent mechanism that only allows the expression of one real value (i.e., $\Theta_i = \mathbb{R}$). (This follows immediately from Cantor (1890): being able to losslessly map between the spaces \mathbb{R}^d and \mathbb{R} .)³*

Thus, it is not the number of real-valued questions that a mechanism can ask that truly characterizes expressiveness, it is how the answers are used!

³Due to the large number of theoretical results in this chapter, proofs of all technical claims are located in an appendix at the end of the chapter.

Another natural way in which mechanisms can differ is in the granularity of their outcome spaces. For example, auction mechanisms that are restricted to allocating certain items together (e.g., blocks of neighboring frequency bands) have coarser outcome spaces than those that can allocate them individually. Some prior work addresses the impact of a mechanism’s *outcome space* on its efficiency. For example, it has been shown that, in private values settings, VCG mechanisms with less coarse outcome spaces always have more efficient dominant-strategy equilibria [74, 104].

In contrast, we are interested in studying the impact of a mechanism’s *expressiveness* on its efficiency. We do this by comparing more and less expressive mechanisms with the *same* outcome space (e.g., fully expressive CAs and multi-item auctions that allow bids on individual items only). In our approach, the outcome space can be unrestricted or restricted; thus the results can be used in conjunction with those stating that larger outcome spaces beget greater efficiency. Furthermore, in many practical applications there is no reason to restrict the outcome space,⁴ but there may be a prohibitive burden on agents if they are asked to express a large amount of information; thus it is limited expressiveness that is the crucial issue.

2.3.1 Impact-based expressiveness

In order to properly differentiate between expressive and inexpressive mechanisms with the same outcome space, we propose to measure the extent to which an agent can impact the outcome that is chosen. We define an *impact vector* to capture the impact of a particular expression by an agent under the different possible types of the other agents. (The subscript $-i$ refers to all the agents other than agent i .)

Definition 6 (impact vector). *An impact vector for agent i is a function, $g_i : T_{-i} \rightarrow \mathcal{O}$. To represent the function as a vector of outcomes, we order the joint types in T_{-i} from 1 to $|T_{-i}|$; then g_i can be represented as $[o_1, o_2, \dots, o_{|T_{-i}|}]$.*

⁴This is the case as long as the mechanism designer’s goal is efficiency, but this is not always the case for revenue maximization, for example. When the designer’s goal is revenue it can be beneficial to restrict the outcome space to induce false competition by, for example, grouping two unrelated products together in an auction.

We say that agent i can *express* an impact vector if there is some pure strategy profile of the other agents such that one of i 's expressions causes each of the outcomes in the impact vector to be chosen by the mechanism.

Definition 7 (express). *Agent i can express an impact vector, g_i , if*

$$\exists h_{-i}, \exists \theta_i, \forall t_{-i}, f(\theta_i, h_{-i}(t_{-i})) = g_i(t_{-i}).$$

We say that agent i can *distinguish* among a set of impact vectors if it can express each of them against the same pure strategy profile of the other agents by changing only its own expression.

Definition 8 (distinguish). *Agent i can distinguish among a set of impact vectors, G_i , if*

$$\exists h_{-i}, \forall g_i \in G_i, \exists \theta_i, \forall t_{-i}, f(\theta_i, h_{-i}(t_{-i})) = g_i(t_{-i}).$$

When this is the case, we say $D_i(G_i)$ is true.

Figure 2.1 illustrates how an agent can distinguish between two different impact vectors against a pure strategy profile of the other agents.




Figure 2.1: By choosing between two expressions, $\theta_i^{(1)}$ and $\theta_i^{(2)}$, agent i can distinguish between the impact vectors $[A, B]$ and $[C, D]$ (enclosed in rectangles). The other agents are playing the pure strategy profile $[\theta_{-i}^{(x)}, \theta_{-i}^{(y)}]$.

Intuitively, more expressive mechanisms allow agents to distinguish among larger sets of impact vectors. Our first expressiveness measure captures this intuition; it measures the number of different impact vectors that an agent can distinguish among. Since this depends on what the others express, we measure the best case for a given agent, where the others submit expressions that maximize the agent’s control. We call this the agent’s *maximum impact dimension*.

Definition 9 (maximum impact dimension). *Agent i has maximum impact dimension d_i if the largest set of impact vectors, G_i^* , that i can distinguish among has size d_i . Formally,*

$$d_i = \max_{G_i} \{|G_i| \mid D_i(G_i)\}.$$

We will show in Section 2.4 that every agent’s maximum impact dimension ties directly to an upper bound on the expected efficiency of the mechanism’s most efficient Nash equilibrium. In particular, the upper bound increases *strictly* monotonically as the maximum impact dimension for any agent i increases from 1 to d_i^* , where d_i^* is the smallest maximum impact dimension needed by the agent in order for the bound to reach full efficiency.

However, the maximum impact dimension also has some drawbacks as a measure. First, it does not capture the way in which an agent’s impact vectors are distributed. For example, it is possible that a mechanism that allows a smaller maximum impact dimension can be designed to let an agent distinguish among a more important (e.g., for efficiency) set of impact vectors. Second, the maximum impact dimension is not well defined in settings where even a single agent has an infinite type space.

2.3.2 Shattering-based expressiveness

We will now discuss a related notion of expressiveness, which we call *outcome shattering*. As we will show later, it has somewhat different uses than the maximum impact dimension.

Outcome shattering is based on a notion called *shattering*, a measure of functional complexity that we have adapted from the field of computational learning theory [27, 147]. In learning theory, a class of binary classification functions⁵ is said to shatter a set of k instances

⁵Binary classification functions are functions that assign each possible input a binary output label of either 0 or 1.

if there is at least one function in the class that assigns each of the possible 2^k dichotomies of labels to the set of instances. Intuitively, a class of functions that can shatter larger sets of instances is more expressive. To illustrate this idea consider the following example taken from Mitchell pp. 215-216 [100].

Example 2. *Consider the class of binary classification functions that assign a 1 to points only if they fall in an interval on the real number line between two constants a and b . Now we can ask whether or not this class of functions has enough expressive power to shatter the set of instances $S = \{3.1, 5.7\}$? Yes, for example the four functions $(1 < x < 2)$, $(1 < x < 4)$, $(4 < x < 7)$ and $(1 < x < 7)$ will assign all possible labels to the instances in S .*

Our adaptation of shattering for mechanisms captures an agent's ability to distinguish among each of the $|\mathcal{O}'|^{T-i}$ impact vectors involving outcomes from a given set, \mathcal{O}' .

Definition 10 (outcome shattering). *A mechanism allows agent i to shatter a set of outcomes, $\mathcal{O}' \subseteq \mathcal{O}$, over a set of joint types for the other agents, T_{-i} , if $D_i(G_i^{\mathcal{O}'})$, where,*

$$G_i^{\mathcal{O}'} = \{g_i \mid g_i = [o_1, o_2, \dots, o_{|T_{-i}|}], o_j \in \mathcal{O}'\} .$$

Example 3. *Suppose the agents other than i have two joint types, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$. If agent i can distinguish among the following set of impact vectors, G_i , by changing only its own expression while the other agents' strategy remains fixed then it can shatter a set of outcomes, $\{A, B, C, D\}$, over the two joint types of the other agents:*

$$G_i = \left\{ \begin{array}{cccc} [A, A], & [B, A], & [C, A], & [D, A], \\ [A, B], & [B, B], & [C, B], & [D, B], \\ [A, C], & [B, C], & [C, C], & [D, C], \\ [A, D], & [B, D], & [C, D], & [D, D] \end{array} \right\}$$

We also use a slightly weaker adaptation of shattering for analyzing the more restricted setting where agents have private values. It captures an agent's ability to cause each of the $\binom{|\mathcal{O}'|+1}{2}$ unordered pairs of outcomes (with replacement) to be chosen for every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., under which of the other agents' types each of the outcomes is chosen). We call this

semi-shattering.⁶

Definition 11 (outcome semi-shattering). *A mechanism allows agent i to semi-shatter a set of outcomes, \mathcal{O}' , over a set of joint types for the other agents, T_{-i} , if i can distinguish among a set of impact vectors, G_i , such that for every pair of joint types $\{x, y \mid x, y \in T_{-i}, x \neq y\}$, and every pair of outcomes, $\{o_1, o_2 \mid o_1, o_2 \in \mathcal{O}', o_1 \neq o_2\}$,*

$$[(\exists g_i \in G_i, g_i(x) = o_1 \wedge g_i(y) = o_2) \wedge \neg(\exists g_i \in G_i, g_i(x) = o_2 \wedge g_i(y) = o_1)] \vee$$

$$[(\exists g_i \in G_i, g_i(x) = o_2 \wedge g_i(y) = o_1) \wedge \neg(\exists g_i \in G_i, g_i(x) = o_1 \wedge g_i(y) = o_2)].$$

The notion of outcome semi-shattering is best illustrated by the following simple examples.

Example 4. *If agent i can distinguish among the following set of impact vectors, G_i , then it can semi-shatter a set of outcomes, $\{A, B, C, D\}$, over two joint types of the other agents (the order of the pairs that are included does not matter, for example $[A, B]$ could be replaced with $[B, A]$):*

$$G_i = \left\{ \begin{array}{l} [A, A], \\ [A, B], [B, B], \\ [A, C], [B, C], [C, C], \\ [A, D], [B, D], [C, D], [D, D] \end{array} \right\}$$

Since semi-shattering is a pairwise notion, it does not always include the entire bottom left half of a sorted matrix, as in the previous example. For example, the following set of impact vectors constitutes semi-shattering a set of three outcomes.

Example 5. *If agent i can distinguish among the following set of impact vectors, G_i , then*

⁶There are many ways to generalize the shattering notion to functions that can return more than two outcomes, c.f. [17]. We have adapted the two most natural ones for our work on expressiveness in mechanism design—in Definitions 10 and 11, respectively. Definition 11 has been slightly altered compared to the version presented at a conference in order to be able to also prove ties to communication complexity.

it can semi-shatter the set of outcomes $\{A, B, C\}$ over three joint types of the other agents:

$$G_i = \left\{ \begin{array}{l} [A, A, A], \\ [A, A, B], \\ [A, A, C], \\ [A, B, B], [B, B, B], \\ [A, C, B], [B, C, B], \\ [A, C, C], [B, C, C], [C, C, C] \end{array} \right\}$$

Notice that every pair of outcomes appears in every pair of slots, in the same order, and at least once, which is exactly the requirement for semi-shattering.

Our second measure of expressiveness is based on the size of the largest outcome space that an agent can shatter or semi-shatter.⁷ It captures the number of outcomes that the mechanism can support full expressiveness over for that agent. We call it the *(semi-)shatterable outcome dimension*.

Definition 12 ((semi-)shatterable outcome dimension). *Agent i has (semi-)shatterable outcome dimension k_i if the largest set of outcomes that i can (semi-)shatter has size k_i .*

The (semi-)shatterable outcome dimension measure addresses both of the concerns with maximum impact dimension that we raised at the end of the previous section. Unlike the maximum impact dimension, which provides no information as to how the distinguishable impact vectors are distributed, the (semi-)shatterable outcome dimension measures the *number of different outcomes* for which an agent has full expressiveness. In addition, it has the advantage that we can rule out the (semi-)shatterability of a set of outcomes by merely ruling out the existence of a *pair* of expressions by the other agents that allows the agent to (semi-)shatter the set.

Observation 1. *Agent i can (semi-)shatter an outcome space \mathcal{O}' only if there exists at least one pair of expressions by the other agents that allows i to (semi-)shatter \mathcal{O}' . (In other*

⁷The measure deals with the size of this space, rather than the specific outcomes it contains, because a designer can always re-label the outcomes in the set to transform it into any other set of the same size.

words, there exists a pair of fixed expressions by the other agents such that i can cause any (unordered) pair of outcomes from \mathcal{O}' to be chosen.)

This observation will allow us to analyze the measure even when agents have infinite type spaces, and may help one operationalize expressiveness for automated mechanism design [45] in the future, since it provides an easy constraint that can be checked to guarantee expressiveness is below a given limit. We use this insight throughout the study of channel-based mechanisms in Section 2.6.

The next two results illustrate the close relationship between the shatterable outcome dimension measures and the maximum impact dimension measure. While the two measures are related, the shatterable outcome dimension can be thought of as a measure of expressiveness breadth.

Proposition 2. *When designing an outcome function, f , increasing a limit on the shatterable or semi-shatterable outcome dimension allowed for a given agent also increases a limit on that agent's maximum impact dimension.*

Proposition 3. *In order to shatter k_i outcomes, agent i must be able to distinguish among at least $|T_{-i}|^{k_i}$ impact vectors.*

Proposition 3 states that the maximum impact dimension necessary for an agent to shatter k outcomes increases geometrically in the number of types of the other agents, which illustrates the relationship between expressiveness and uncertainty. As uncertainty goes up (the number of types that the other agents have can be thought of as a support-based measure of uncertainty), more expressiveness is needed to shatter a given set of outcomes.

2.3.3 Uses of the expressiveness measures

The expressiveness measures introduced above enable us to understand mechanisms from a new perspective. Because the measures are so new, we undoubtedly fail to see all of their possible uses at this time, however we already see several.

First, we can measure the expressiveness of an existing mechanism, and thereby bound how well the mechanism can do in terms of a designer's objective. For example, in the next

section, we show how our expressiveness measures directly relate to an upper bound on the efficiency of any mechanism.

Second, one may be able to use the expressiveness measures in designing new mechanisms. For example, if there are some constraints on what—and how much—information the agents can submit to the mechanism (e.g., in a CA, allowing bids on packages of no more than k items), then our measures can be used to design the most expressive mechanism subject to those constraints. This, in turn, hopefully maximizes the mechanism designer’s objective subject to the constraints. For example, our results presented in the next section imply that this approach can be used to yield the most efficient possible Bayes-Nash equilibrium in any private values setting.

We can also ask which of the expressiveness measures—maximum impact dimension, shatterable outcome dimension, or semi-shatterable outcome dimension—is most appropriate under different settings and for different purposes. If the designer knows which impact vectors are (most) important, then the maximum impact dimension is the measure of choice. If, instead, the designer knows which outcomes are (most) important but not which impact vectors are (most) important, then the other two measures can be used to make sure that the agents have full expressiveness over those outcomes. As we will show in Section 2.4, in private values settings the appropriate measure is semi-shatterable outcome dimension (for one, semi-shatterability is enough to guarantee that lack of expressiveness will not limit the mechanism’s efficiency at all), and in interdependent values settings the appropriate measure is shatterable outcome dimension. Also, we will show that less than full (semi-)shatterability necessarily leads to arbitrary inefficiency under some preference distributions.

Another use of the semi-shatterable outcome dimension is to analyze a broad subclass of mechanisms which we will call channel based. This will be discussed in Section 2.6.

2.4 Expressiveness and efficiency

Perhaps the most important property of our domain-independent measures of expressiveness is how they relate to the efficiency of the mechanism’s outcome. In this section, we will discuss a cooperative upper bound on the expected efficiency of a mechanism’s *most efficient* equilibrium that is tied directly to the expressiveness of an optimally designed mechanism

and can always be implemented by a budget-balanced mechanism in Bayes-Nash equilibrium (in private values settings).⁸

The bound measures the efficiency of the outcome function under the optimistic assumption that the agents play strategies which, taken together, attempt to maximize expected efficiency. Studying this bound allows us to sidestep two of the major roadblocks faced by many prior attempts at analyzing the relationship between expressiveness and efficiency: 1) we do not have to solve for any of the mechanism's equilibria (attempts at doing this have proved difficult for many expressive and inexpressive mechanisms [103, 107, 119, 139, 155, 159]) and 2) since it bounds the *most efficient* equilibrium, it can be used to study mechanisms with multiple—or an infinite number of—equilibria, e.g., first price CAs [24]. This allows us to avoid the difficulty involved in calculating equilibrium strategies. It also implies that we can restrict our analysis to *pure* strategies because a pure strategy always exists that achieves at least as much expected efficiency as any mixed strategy.

Proposition 4. *The following quantity, $E[\mathcal{E}(f)]^+$, is an upper bound on the expected efficiency of the most efficient mixed-strategy profile under any mechanism with outcome function f ,*

$$(2.2) \quad E[\mathcal{E}(f)]^+ = \max_{\hat{h}(\cdot)} \int_{t^n} P(t^n) W(f(\hat{h}(t^n)), t^n).$$

The bound holds for mixed strategies, but the maximum in the equation need only be taken over the space of pure-strategy profiles, $\hat{h}(\cdot)$.

To see how this bound is tied to our notions of expressiveness, consider calculating it from the fixed perspective of a particular agent i . Based on the assumption behind the bound, the other agents will choose whatever pure strategies are best for maximizing expected efficiency. Thus, from agent i 's perspective, the maximization above amounts to finding the set of expressible impact vectors that lead to the highest expected efficiency.

⁸The upper bound we derive represents a cooperative equilibrium that could be used to bound the value of any objective that depends only on the agents' types and the outcome chosen by the mechanism. By extension, all of our subsequent theory (except for the implementability of the bound discussed in Section 2.4.4) also applies for any such objective.

2.4.1 Conditions for full efficiency

There is an impact vector for each of agent i 's types that represents the vector of the most efficient outcomes when it is matched with each of the joint types of the other agents. In order to achieve full efficiency, agent i must be able to distinguish among all of these vectors. We call a set of such impact vectors a *fully efficient set*. Such a set must be distinguishable by each agent for the bound to reach full efficiency.

Definition 13 (fully efficient set). *A set of unique impact vectors, G_i^* , for agent i is a fully efficient set if it contains an impact vector corresponding to the vector of efficient outcomes when each of agent i 's types is matched with all of the non-zero probability types of the other agents. Formally, G_i^* is a fully efficient set if*

$$\forall t_i, \exists g_i \in G_i^*, \forall \{t_{-i} \mid P(t_i, t_{-i}) > 0\}, W(g_i(t_{-i}), [t_i, t_{-i}]) = \max_{o \in \mathcal{O}} W(o, [t_i, t_{-i}]).$$

Our next two results demonstrate that in order to achieve full efficiency, an outcome function must allow each agent to distinguish among one of its fully efficient sets.

Proposition 5. *$E[\mathcal{E}(f)]^+$ reaches full expected efficiency if and only if each agent can distinguish among the impact vectors in at least one of its fully efficient sets.*

Proposition 6. *If any agent can distinguish among each of the impact vectors in at least one of its fully efficient sets, then each other agent can also distinguish among each of the impact vectors in at least one of its fully efficient sets.*

In full information settings, whereupon learning its own type an agent knows the types of the other agents for sure, the agent is guaranteed to have a fully efficient set of size $\leq |\mathcal{O}|$. (This is slightly more general than assuming the agent has perfect information about the types of the other agents *a priori*, since it need only have this information once its own type is revealed.)

Proposition 7. *Let G_i^* be agent i 's smallest fully efficient set,*

$$(\forall t_i, \exists t_{-i} \mid P(t_i, t_{-i}) = 1) \Rightarrow |G_i^*| \leq |\mathcal{O}|.$$

Corollary 1. *If agent i has full information then there exists an outcome function for which the upper bound reaches full efficiency while limiting i to maximum impact dimension $d_i \leq |\mathcal{O}|$.*

The implication of this result is that perfect information about the other agents' types essentially eliminates the need for expressiveness. Thus, in prior research showing that in certain settings even quite inexpressive mechanisms yield full efficiency in ex-post Nash equilibrium (e.g., [1]), the assumption that the agents know each other's types is likely essential.

2.4.2 The efficiency bound increases strictly with expressiveness

The following results demonstrate that a mechanism designer can *strictly* increase the upper bound on expected efficiency by giving any agent more expressiveness (until the bound reaches full efficiency). The result applies to the outcome function that maximizes the bound subject to the constraint that agent i 's expressiveness be less than or equal to a particular level. The bound attained by such an optimal outcome function is also an upper bound for any outcome function at that expressiveness level or lower.

Theorem 1. *For any distribution over agent preferences, the upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, for the best outcome function limiting agent i to maximum impact dimension d_i increases strictly monotonically as d_i goes from 1 to d_i^* (where d_i^* is the size of agent i 's smallest fully efficient set).*

From Proposition 2, we know that any increase in allowable shatterable or semi-shatterable outcome dimension implies an increase in allowable maximum impact dimension; thus Theorem 1 implies that strict monotonicity holds for these measures as well.

Corollary 2. *The upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$, of the best outcome function limiting agent i 's expressiveness to (semi-)shatterable outcome dimension k_i increases strictly monotonically as k_i goes from 1 to k_i^* (where k_i^* is the (semi-)shatterable outcome dimension necessary for the bound to reach full efficiency).*

2.4.3 Inadequate expressiveness can lead to arbitrarily low efficiency for some preference distributions

The next three lemmas provide the foundation for our second main theorem regarding the efficiency bound. They demonstrate that in *any* setting there are distributions over agent

preferences under which an increase in allowed expressiveness leads to an arbitrary improvement in the upper bound on expected efficiency. We prove that the arbitrary increase is possible by constructing an example under which it is inevitable. (We try to keep these constructions as general as possible: we allow for any number of outcomes, any number of agents, and any number of types.)

Lemma 1. *For any agent i in an interdependent values setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents), there exist preference distributions under which $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i 's maximum impact dimension to d_i (where $2 \leq d_i \leq |\mathcal{O}|^{T-i}$) is arbitrarily larger than that of any outcome function limiting i 's maximum impact dimension to $d_i - 1$.*

The next lemma deals with the arbitrary improvement that can be achieved by allowing an agent to shatter a single additional outcome. Here we distinguish between an increase in shatterable outcome dimension, for interdependent values settings, and semi-shatterable outcome dimension, for private values settings. As we will see, in a private values setting there is no need to allow full shattering to achieve full efficiency.

Lemma 2. *For any agent i in any setting (with any number of outcomes, any number of other agents, and any number of joint types for those agents), there exist preference distributions under which $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i 's expressiveness to*

- shatterable outcome dimension k_i for interdependent values settings, or
- semi-shatterable outcome dimension k_i for private values settings

(where $2 \leq k_i \leq |\mathcal{O}|$) is arbitrarily larger than that of any outcome function that limits i 's expressiveness to $k_i - 1$.

Private values settings place restrictions on agents' utility functions and, therefore, on the efficiency-maximizing outcomes under different types. We will use the following lemma to show that in such settings allowing the agents to semi-shatter the outcomes is sufficient for maximizing the efficiency bound. The lemma proves that the most efficient *order* for two outcomes under any pair of opposing types must be the same for all of agent i 's types.

Lemma 3. *In any private values setting, for any agent i , any pair of outcomes, o_1 and o_2 , and any pair of types for the other agents, $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, if there exists some type of agent i , t_i , where it is strictly more efficient for o_1 to be chosen under type $t_{-i}^{(1)}$ and o_2 to be chosen under type $t_{-i}^{(2)}$ than vice-versa (i.e., o_1 for $t_{-i}^{(2)}$ and o_2 for $t_{-i}^{(1)}$), then it cannot be more efficient for the outcomes to be chosen in the other order for any type of agent i .*

We conclude this section with a result that integrates the three lemmas above. The theorem adds the fact that an arbitrary loss in efficiency can *only* happen if the shatterable (for interdependent values) or semi-shatterable (for private values) outcome dimension is less than the number of outcomes in the mechanism. Thus, these dimensions can be used to provide a guarantee that a mechanism has enough expressiveness to avoid arbitrary expected efficiency loss under any possible preference distribution.

Theorem 2. *For any setting, there exists a distribution over agent preferences such that $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i to*

- shatterable outcome dimension, $k_i < |\mathcal{O}|$, in an interdependent values setting, or
- semi-shatterable outcome dimension, $k_i < |\mathcal{O}|$, in a private values setting

is arbitrarily less than that of the best outcome function limiting agent i to (semi-)shatterable outcome dimension $k_i + 1$.

Since we have identified a gap in an upper bound on expected efficiency, our results in this section demonstrate that any mechanism that does not allow any agent to shatter (in interdependent values settings) or semi-shatter (in private values settings) its entire outcome space will be arbitrarily inefficient under some preference distribution.

2.4.4 Bayes-Nash implementation of the upper bound is always possible in private values settings

In addition to the results above, we find that the upper bound on expected efficiency can be implemented in Bayes-Nash equilibrium for any outcome function, in any private values⁹

⁹Implementing efficient allocations in Bayes-Nash equilibrium for interdependent values settings is impossible even with full expressiveness [83]. The difficulty stems from the need for the mechanism designer to know the beliefs of the agents about *each others'* private information.

setting, as long as the agents have quasi-linear utility functions. Quasi-linearity means that the agent's utility functions are linear in money or some commonly agreed upon currency. Formally, a quasi-linear utility function for agent i takes the form $u_i = v_i - \pi_i$, where v_i is the agent's valuation for the outcome chosen by the mechanism and π_i is the payment from the agent to the mechanism. This is the first point in the paper where we assume quasi-linearity: all the results so far apply with and without that restriction.

Theorem 3. *For any private values setting with quasi-linear preferences and any outcome function, f , there exists a class of payment functions that achieve the upper bound on efficiency, $E[\mathcal{E}(f)]^+$, in a pure-strategy Bayes-Nash equilibrium.*

This implementability of the upper bound implies that, for private values settings, we can recast all of our earlier results that relate expressiveness to the bound as relating expressiveness to the efficiency of the most efficient implementable Bayes-Nash equilibrium.

Individual rationality and budget balance

In this section, we will discuss individual rationality and budget balance, and how they are related to expressiveness. First, in Bayes-Nash equilibrium, we can always get *strong* budget balance (i.e., the total payments to and from all agents are equal), and we can get *ex ante* individual rationality (i.e., it is always in an agent's best interest to participate in the mechanism prior to learning its own type) as long as agent valuations for outcomes satisfy the following criterion, which is generally satisfied in most commonly studied settings (e.g., it is satisfied in all auction settings).

Definition 14 (Non-negative externality criterion). *A preference distribution satisfies the non-negative externality criterion for a given outcome function, f , if the expected welfare of every group of $n - 1$ agents is non-negative when agents play welfare-maximizing strategies, i.e., $\forall i, E_t[W_{-i}(f(h^*(t)), t_{-i})] \geq 0$.*

Proposition 8. *There exists at least one payment function in the class of Theorem 3 that is strongly budget balanced and, as long as the preference distribution satisfies the non-negative externality criterion, ex ante individually rational.*

These payment functions are derived much like in the expected-form Groves mechanism which is due to d'Aspremont and Gerard-Varet [58] and Arrow [8] (called the dAGVA mechanism). However, as implied by the Myerson-Satterthwaite impossibility theorem [102] for fully expressive mechanisms, there may not exist a payment function in this class that is *ex interim* individually rational (i.e., it may not be in an agent's best interest to participate once the agent knows its own type). Additionally, there exist settings, such as the one described in the following example, where the fully expressive dAGVA mechanism is ex-interim individually rational but a limited-expressiveness variant is not (see Chapter 2 in Parkes' Ph.D. dissertation [112] for a thorough discussion of dAGVA).

Example 6. *Consider an auction for one item run using the dAGVA mechanism. Assume there are two bidders with valuations for the item drawn from the uniform distribution over $[0, 1]$ and one auctioneer with zero valuation for the item. Let $\hat{\theta}$ represent the bidders' bids, assuming they follow the Bayes-Nash equilibrium and report their valuations truthfully, and let $f_i(\theta)$ be an indicator function that returns one if bidder i wins the item and zero otherwise. The following reasoning demonstrates that the fully expressive dAGVA mechanism is ex interim individually rational for this setting.*

First, we can calculate the dAGVA payment for one of the bidders, i , for a given set of bids (the payment to the auctioneer is the sum of the payments from the bidders). We begin with the general formula for the dAGVA payment function in a direct-revelation mechanism and then instantiate it for a bidder in this example.

$$\begin{aligned} \pi_i(\hat{\theta}) &= -E_{\theta_{-i}}[W_{-i}(f(\hat{\theta}_i, \theta_{-i}), \theta_{-i})] + \frac{1}{n-1} \left(\sum_{j \neq i} E_{\theta_{-j}} [W_{-j}(f(\hat{\theta}_j, \theta_{-j}), \theta_{-j})] \right) \\ &= -E_{\theta_{-i}}[\theta_{-i} f_{-i}(\hat{\theta}_i, \theta_{-i})] + \frac{1}{2} \left(E_{\theta_i}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] + E_{\theta}[\max(\theta_i, \theta_{-i})] \right) \\ &= \frac{1}{12} + \frac{\hat{\theta}_i^2}{2} - \frac{\hat{\theta}_{-i}^2}{4} \end{aligned}$$

Next, we can calculate bidder i 's expected utility when it draws type θ_i .

$$\begin{aligned} E[u_i|\theta_i] &= E_{\hat{\theta}_{-i}}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] - E_{\hat{\theta}_{-i}} \left[\pi_i(\theta_i, \hat{\theta}_{-i}) \right] \\ &= E_{\hat{\theta}_{-i}}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] - E_{\hat{\theta}_{-i}} \left[\frac{1}{12} + \frac{\theta_i^2}{2} - \frac{\hat{\theta}_{-i}^2}{4} \right] \\ &= \frac{\theta_i^2}{2} \end{aligned}$$

Since θ_i can never be less than zero, $E[u_i|\theta_i]$ can never be negative in this setting.

However, if we consider bidder i 's expected utility under the best outcome function with less than full expressiveness, we find this is not necessarily true. Say we limit the expressiveness to a maximum impact dimension of one, which entails creating a deterministic mechanism that always chooses the same outcome. Now, the best limited-expressiveness outcome function for this setting always allocates the item to the same bidder regardless of the bids. Call that bidder i . Under these assumptions, bidder i 's dAGVA payment and expected utility are

$$\begin{aligned} \pi_i(\hat{\theta}) &= -E_{\theta_{-i}}[\theta_{-i}f_{-i}(\hat{\theta}_i, \theta_{-i})] + \frac{1}{2} \left(E_{\theta_i}[\theta_i f_i(\theta_i, \hat{\theta}_{-i})] + E_{\theta}[\theta_i] \right) \\ &= E_{\theta_i}[\theta_i] \\ E[u_i|\theta_i] &= \theta_i - E_{\theta_i}[\theta_i] \end{aligned}$$

Thus, bidder i 's expected utility is negative whenever the valuation it draws is less than its expected valuation.

Impossibility of dominant strategy implementation

While Theorem 3 shows it is always possible to implement the upper bound for private values settings in Bayes-Nash equilibrium, we show below that there exist private values settings for which dominant strategy implementation is impossible without full expressiveness. In other words, it is known that with full expressiveness there is no difference between what is possible in dominant strategies and Bayes-Nash equilibrium (except for issues of individual rationality and budget balance), but we show that with less than fully expressive mechanisms there is a fundamental difference in the power of the two solution concepts.

Theorem 4. *There exist private values settings with quasi-linear preferences where the outcome function that maximizes the upper bound on efficiency, $E[\mathcal{E}(f)]^+$, while limiting agent i to a maximum impact dimension $d_i < d_i^*$ (d_i^* is the size of agent i 's smallest fully efficient set), cannot be implemented in dominant strategies.*

The reason for this impossibility is that there exist settings where the best limited-expressiveness outcome function is not guaranteed to satisfy the *weak-monotonicity* property, a condition which has been shown to be necessary for dominant strategy implementation [25].

This property requires that the outcome function react properly to relative changes in an agent's reported preferences for any two outcomes.

2.5 Expressiveness and communication complexity

In this section, we consider the relationship between our notions of expressiveness and more traditional notions of communication complexity. Our expressiveness measures quantify how the mechanism uses information, while communication complexity measures how much information has to be communicated (by the agents) to compute it. Although these notions do not measure exactly the same thing, they are closely related. In this section we will begin to formalize this relationship.

One measure of an outcome function's communication complexity for agent i is the size of its expression space, $|\Theta_i|$. As we will show, this determines an upper bound on the amount of information communicated by the agent under any communication procedure that computes the outcome function.

In relating expressiveness to the number of expressions needed for each agent, we consider whether or not a given outcome function can be *emulated* by an outcome function with fewer expressions (essentially losslessly compressed). We say an outcome function, f' , emulates another outcome function, f , if there exists a one-to-many mapping for each agent from expressions in f to expressions in f' such that the outcomes chosen by f' under the mapping are the same as those chosen by f under the original expressions.

Definition 15 (emulate). *An outcome function, f' , emulates another outcome function, f , if there exists a function, q_i , for each agent, i , that maps from i 's expression space under f to i 's expression space under f' , such that*

$$\forall i, \forall \theta_i, \forall \theta_{-i}, f(\theta_i, \theta_{-i}) = f'(q_i(\theta_i), q_{-i}(\theta_{-i})).$$

For a given outcome function, f , each agent's maximum impact dimension provides a lower bound on the the number of expressions needed for that agent by any outcome function that emulates f .

Proposition 9. *It is impossible to emulate an outcome function, f , with an outcome function that provides any agent with less expressions than its maximum impact dimension under f .*

Furthermore, for any outcome function that belongs to the widely studied class of direct-revelation mechanisms, an agent's maximum impact dimension is exactly the number of expressions used by the best emulator of the outcome function (i.e., the outcome function that emulates it while minimizing the number of expressions).

Lemma 4. *Under a direct-revelation outcome function, each agent i 's impact dimension is maximized when the agents other than i report their types truthfully, (i.e., $h_{-i}(t_{-i}) = t_{-i}$ is the strategy that maximizes i 's impact dimension).*

Proposition 10. *Any direct-revelation outcome function, f , can be emulated by another outcome function, f' , that provides each agent, i , with exactly d_i expressions, where d_i is agent i 's maximum impact dimension under f .*

Given this relationship between our expressiveness measure and the number of expressions needed by any agent, we have the following two Corollaries related to the upper bound on expected efficiency, $E[\mathcal{E}(f)]^+$. Corollary 3 states that increasing a limit on the number of expressions given to an agent strictly increases the bound. Corollary 4 states that some distributions require an agent to have a number of expressions that is exponential in the number of types of the other agents to avoid being arbitrarily less than fully efficient.

Corollary 3. *For any setting and any distribution over agent preferences, $E[\mathcal{E}(f)]^+$ for the best outcome function limiting agent i to d_i expressions increases strictly monotonically as d_i goes from 1 to d_i^* , where d_i^* is the size of agent i 's smallest fully efficient set.*

Corollary 4. *There exists settings and distributions over agent preferences such that the upper bound on expected efficiency for the best outcome function limiting agent i to less than $|T_{-i}|^{|O|}$ expressions is arbitrarily less than that of the best outcome function.*

While the reasoning above provides an upper bound on an outcome function's communication complexity, it does not account for the possibility of designing clever elicitation protocols, such as protocols that iteratively ask different agents different questions (cf. [126]). To address this, we will also relate our notion of expressiveness to a lower bound on communication complexity. The lower bound is derived by considering the execution of the outcome function as a two-party communication problem, where agent i holds one piece of information (its intended expression) and the agents other than i hold another (their intended joint

expression). From this perspective, we can study the outcome function using Yao’s model of communication complexity [158], as in Nisan and Segal’s seminal work on communication complexity in mechanism design [105].

Yao’s model considers the computation of a pre-specified function based on the information held by the agents. It is typical, when using this model, to think of the function being computed as a matrix where the rows represent the possible inputs to the function from one agent, the columns represent the possible inputs from the other agent, and each cell contains the value of the function under the inputs corresponding to its row and column. For a given outcome function, f , and agent, i , we can construct such an input matrix from i ’s perspective by placing its expressions along the rows and the joint expressions of the other agents along the columns. The cells of the matrix contain the outcome chosen by the outcome function under the corresponding expressions. (Thus, the rows correspond to agent i ’s possible impact vectors.)

It has been shown that any communication protocol that computes f must involve at least one message for each of the *monochromatic rectangles* (i.e., contiguous rectangles of expressions that result in the same outcome being chosen) in some partitioning of f ’s input matrix [87]. The following result shows how our notion of semi-shattering is related to the number of monochromatic rectangles needed in any partitioning of f . Specifically, any set of types for the agents other than i over which agent i can semi-shatter a pair of outcomes leads to a corresponding set of expression pairs that cannot be in the same monochromatic rectangle for either outcome.

Lemma 5. *Let T_{-i} be a set of joint types for the agents other than i over which agent i can semi-shatter a pair of outcomes, A and B , under some outcome function, f . There exists a set of $|T_{-i}| - 1$ pairs of expressions that cannot be in the same A - or B -monochromatic rectangle of f .*

This leads directly to a lower bound on the number of monochromatic rectangles needed by any partitioning of an outcome function’s input matrix and, consequently, a lower bound on the number of messages needed by any communication protocol that computes it.

Theorem 5. *For any outcome function, f , agent, i , and outcome, o , let T_{-i}^o denote the largest set of types over which i can semi-shatter a pair of outcomes containing o . Also, let*

d_i be i 's maximum impact dimension under f . The number of monochromatic rectangles, R , needed by any partitioning of the input matrix of f , and the number of messages needed by the best communication protocol that computes f , $M(f)$, satisfy the following inequality,

$$(2.3) \quad \min_i d_i \geq M(f) \geq R \geq \max_i \sum_{o \in O} |T_{-i}^o| - 1.$$

This result bounds the number of bits needed by any discrete communication protocol that computes f , since a function that requires $M(f)$ messages to compute must communicate at least $\log_2(M(f))$ bits (i.e., the depth of a binary tree with $M(f)$ leaves). Our bound is also consistent with earlier results showing that combinatorial allocation mechanisms can require the communication of a number of bits that is exponential in the number of items [105], since the number of types an agent has in a combinatorial allocation setting is typically doubly-exponential in the number of items: if there are m items, and an agent has k possible values for each bundle, then the agent has k^{2^m} types. Thus, according to Theorem 5, a combinatorial allocation mechanism that allows an agent to semi-shatter even a single pair of outcomes over the other agents' entire type space would require at least $\log_2(|T_{-i}| - 1)$ bits, which is on the order of 2^m bits.

2.6 Expressiveness in channel-based mechanisms

We will now instantiate our theory of expressiveness for an important class of mechanisms, which we call *channel based*. Channel-based mechanisms are defined as follows (a small example is also presented in Figure 2.2).

Definition 16 (channel-based mechanism). *Each outcome is assigned a set of channels potentially coming from a number of different agents (e.g., outcome A may be assigned channels x_1 and y_1 from Agent 1 and x_2 from Agent 2). Each agent, simultaneously with the other agents, reports real values on each of its channels to the mechanism. The number of channels assigned to each agent, i , is denoted k_i . The mechanism chooses the outcome whose channels from all agents have the largest sum.¹⁰ Formally, a channel-based mechanism has*

¹⁰We assume that ties are broken consistently according to some strict ordering on the outcomes. This prevents an agent from using the mechanism's tie breaking behavior as artificial expressiveness.

the following properties.

- The expression space of agent i is a vector of real numbers with dimension c_i , (i.e., $\Theta_i \equiv \mathbb{R}^{k_i}$). Each dimension is called a channel.
- For each agent i there is a set of channels associated with each outcome o , S_i^o , such that the mechanism's outcome function chooses the outcome with associated channels that have the greatest reported sum:


$$f(\theta) = \arg \max_{O \in \mathcal{O}} \sum_i \sum_{j \in S_i^O} \theta_{ij}.$$

Many different mechanisms for trading goods, information, and services, such as combinatorial allocation mechanisms, CAs, exchanges, and multi-attribute mechanisms can be cast as channel-based mechanisms. (This class is even more general than CAs because it can model settings where agents care about how the items that they do not win get allocated across the other agents, e.g., an advertisement auction where agents care about which slots are assigned to competitors.)


A natural measure of expressiveness in channel-based mechanisms is the number of channels allowed. For CAs, it is able to capture the difference between fully expressive CAs, multi-item auctions that allow bids on individual items only (Fig. 2.2), and an entire spectrum in between. In fact, it generalizes a classic measure of expressiveness in CAs called k -wise dependence [51]. First, we will demonstrate that our domain-independent expressiveness measures relate appropriately to the number of channels allowed in a channel-based mechanism. The following result shows that as the number of allowed channels for an agent increases, the agent's expressiveness in the most expressive channel-based mechanism *strictly* increases as well (until full expressiveness is reached). In particular, each time an agent is given an additional channel it is possible to design a channel-based mechanism that allows the agent to semi-shatter over at least one additional outcome.

Proposition 11. *For any agent i , its semi-shatterable outcome dimension, k_i , in the most expressive channel-based mechanism strictly increases (until $k_i = |\mathcal{O}|$) as the number of channels assigned to the agent increases.*

It is interesting to note that, while adding a new channel can result in an increase in expressiveness, it is also possible to add channels that do not lead to increases in expressiveness



Fully expressive combinatorial auction.



Auction that only allows bids on items.

Figure 2.2: Channel-based representations of two auctions. The items auctioned are an apple (a) and an orange (o). The channels for each agent i are denoted x_i , y_i , and z_i . The possible allocations are A, B, C, and D. In each one, the items that Agent 1 gets are in the first braces, and the items Agent 2 gets are in the second braces.

(e.g., by connecting them to the wrong outcomes). This is consistent with our discussion in Section 2.5, where we showed that increases in expressiveness necessitate increases in message space size, but increasing the size of the message space does not always lead to an increase in expressiveness. Rather, it depends on how the resulting mechanism is wired.

Based on Theorems 1 and 2, we know that an increase in expressiveness will always yield an increase in our efficiency bound and can lead to an arbitrarily large increase, even in private values settings.

Corollary 5. *For any setting, the upper bound on expected efficiency of the best channel-based mechanism that allows c_i channels for agent i is strictly greater than, and can be arbitrarily larger than, that of the best mechanism that allows agent i to have c'_i channels, where $c_i < c'_i \leq c_i^*$ and c_i^* is the number of channels needed for full efficiency.*

However, if an agent has full information it only needs a logarithmic number of channels to bring the bound to full efficiency. (This also happens to be the number of channels in any multi-item auction that allows item bids only.)

Proposition 12. *If agent i has full information about the other agents, in a channel-based mechanism it needs only $\lceil \log_2(|\mathcal{O}|) \rceil$ channels to shatter the entire outcome space.*

On the other hand, an agent with less than full information cannot *fully* shatter any set of two or more outcomes in a channel-based mechanism.

Proposition 13. *No channel-based mechanism allows any agent to shatter any set of two or more outcomes when the other agents have two or more types.*

Since channel-based mechanisms do not allow full shattering, our results from the previous section imply that in some interdependent values settings any channel-based mechanism, even one that emulates the VCG mechanism, will be arbitrarily inefficient. However, these mechanisms are typically studied in private values settings where (as demonstrated by Lemma 3) semi-shattering is more important than full shattering for efficiency. (That such mechanisms cannot always get full efficiency in interdependent values settings is already known [83].)

Corollary 6. *In any interdependent values setting, there exist preference distributions for which any channel-based mechanism (even one that emulates the VCG mechanism) results in arbitrarily less than full expected efficiency.*

However, full efficiency can be achieved in any private values setting—despite agent uncertainty—by a channel-based mechanism with $|\mathcal{O}| - 1$ channels per agent that emulates the VCG mechanism.

Proposition 14. *A channel-based mechanism can emulate the VCG mechanism if and only if it provides each agent with at least $|\mathcal{O}| - 1$ channels.*

Our next two results deal with a configuration of channels that prevents an agent from being able to semi-shatter a set containing two particular pairs of outcomes.

We will first present a lemma regarding an implication based on set algebra that will be used to prove the second lemma.

Lemma 6. *For any sets, A , B , C , and D , the following bi-directional implication holds,*

$$(A \setminus C = D \setminus B) \text{ and } (C \setminus A = B \setminus D) \quad \Leftrightarrow \quad (A \setminus D = C \setminus B) \text{ and } (D \setminus A = B \setminus C).$$

Lemma 7. *Consider a set of outcomes, $\{A, B, C, D\}$, connected to different sets of channels for agent i , $\{S_i^A, S_i^B, S_i^C, S_i^D\}$, respectively. Agent i cannot semi-shatter any set of outcomes containing both pairs $\{A, B\}$ and $\{C, D\}$ (i.e., there is no fixed pair of expressions by the other agents allowing i to cause the mechanism to select A and B with one expression, and C and D with another) if,*

$$(S_i^A \setminus S_i^C = S_i^D \setminus S_i^B) \quad \text{and} \quad (S_i^C \setminus S_i^A = S_i^B \setminus S_i^D).$$

An illustration of the channel configuration discussed in Lemma 7 is shown in Figure 2.3. This configuration generalizes one that appears in the channel-based representation of a CA where bids are allowed on items only. In fact, it is present in any combinatorial allocation mechanism whenever it is assumed that an agent’s bid for a bundle is the sum of its bid on two other non-overlapping bundles (e.g., sub-bundles that compose the full bundle). This is true even if the bids on the sub-bundles are complex themselves (i.e., assumed to be the sum of bids on other bundles).




Figure 2.3: An agent controlling non-overlapping sets of channels S_1 , S_2 , and S_3 can semi-shatter a pair of opponent profiles over A and B or C and D but not both.

Based on this insight, we can prove that for any combinatorial allocation mechanism where an agent's bid on any bundle is the sum of its bid on two other non-overlapping bundles, there exists a preference distribution satisfying free disposal (i.e., an agent's valuation for a bundle of items is greater than or equal to its valuation for any sub-bundle) where the mechanism cannot achieve expected efficiency within 5% of the maximum. While 5% may seem like a relatively small gap, it can be arbitrarily large in absolute terms. Furthermore, it is ten times larger than the expected efficiency gap found by Nisan and Segal [105] in their prior work on communication complexity in combinatorial allocation mechanisms. Their result pertains to mechanisms that communicate less than an exponential number of bits and involves a single prior over preferences. Our result pertains to limited expressiveness and potentially uses a different prior for each mechanism.

Theorem 6. *Consider a combinatorial allocation mechanism, M , which can be represented as a channel-based mechanism that treats agent i 's bid on any bundle Q to be the sum of its bids on some two other non-overlapping sub-bundles, q_1 and q_2 . There exists a distribution over agent valuations, that satisfies the private values and free disposal assumptions, such that M cannot achieve expected efficiency within 5% of the maximum possible for the setting.*

The setting that is used to prove Theorem 6 provides some insight into the circumstances under which limited expressiveness can be particularly problematic. It involves two agents who care only about the items with limited expressiveness (i.e., the items in q_1 and q_2). Each of the agents has two equally probable types: a *complementarity* type and a *substitutability* type. Under the complementarity type, the agent only derives utility from winning the super-bundle (i.e., the bundle containing both q_1 and q_2). Under the substitutability type, the agent derives no additional utility from winning more than one sub-bundle (either q_1 or q_2). In other words, we find that expressiveness is important for combinatorial allocation mechanisms when agents may have either complementarity or substitutability for the same items.

2.7 Conclusions and future research

A recent trend in (electronic) commerce is a demand for higher levels of expressiveness in the mechanisms that mediate interactions such as the allocation of resources, matching of peers, or elicitation of opinions. In this paper we provided the first general model of expressiveness for mechanisms. Our model included a new expressiveness measure, maximum impact dimension, that captures the number of different ways an agent can impact the outcome of a mechanism. We also introduced two related measures of expressiveness based on the concept of shattering from computational learning theory.

We then described perhaps the most important property of our domain-independent expressiveness notions: how they relate to the efficiency of the mechanism's outcome. We derived an upper bound on the expected efficiency of a mechanism's most efficient equilibrium that depends only on the extent to which agents can impact the mechanism's outcome. This bound enables us to study the relationship between expressiveness and efficiency by avoiding two major classic hurdles: 1) the bound can be analyzed without having to solve for an equilibrium of the mechanism, and 2) the bound applies to the most efficient equilibrium so it can be used to analyze mechanisms with multiple (or an infinite number of) equilibria. We proved that this bound increases *strictly* monotonically for the best mechanism that can be designed as the limit on any agent's expressiveness increases (until the bound reaches full efficiency). In addition, we proved that a small increase in expressiveness can lead to arbitrarily large increases in the efficiency bound, depending on the prior over

agents' preferences. We ended the discussion with proof that the bound is tight in private values settings: it is always possible to build a strongly budget balanced payment function that achieves the efficiency of the bound in Bayes-Nash equilibrium (with *ex ante* but not necessarily *ex interim* individual rationality). This implies that for any private values setting, the expected efficiency of the best Bayes-Nash equilibrium increases strictly as more expressiveness is allowed. However, we showed that unlike with full expressiveness, implementing the bound is not always possible in dominant strategies with less than full expressiveness. Additionally, the efficiency of the bound may not be achieved by a mechanism if its payment function is not properly designed to incentivize it. Still, these results provide a significant step forward in our understanding of the relationship between expressiveness and efficiency.

Next, we explored the relationship between our expressiveness measures and communication complexity. We showed that the expressiveness measures can be used to derive both upper and lower bounds on the number of bits used by the best communication protocol for running any mechanism.

Finally, we instantiated our model of expressiveness for a class of mechanisms, called channel based. This class involves mechanisms that take expressions of value through channels from agents to outcomes, and select the outcome with the largest sum. Many mechanisms for trading goods, information, and services—such as combinatorial auctions, exchanges, and multi-attribute auctions—can be cast as channel-based mechanisms. We showed that our domain-independent measures of expressiveness appropriately relate to a natural notion of expressiveness in channel-based mechanisms, the number of channels allowed (which already generalizes a traditional measure of expressiveness in CAs called *k*-wise dependence [51]). Using our general measures of expressiveness and the results on how they relate to efficiency, we proved that in channel-based mechanisms 1) increasing expressiveness by allowing an additional channel leads to an increase in the upper bound on expected efficiency for the mechanism, and 2) under some preference distributions this leads to an arbitrarily large increase in the bound. We also used our theoretical framework to prove that for any (channel-based) multi-item allocation mechanism that does not allow rich combinatorial bids, there exist distributions over agent preferences that satisfy the free disposal condition for which the mechanism cannot achieve 95% of optimal efficiency. This inefficiency is ten times larger than a related expected efficiency gap found by Nisan and Segal [105] in their prior work on communication complexity in combinatorial allocation mechanisms.

Our work in this chapter opens up several opportunities for future research. First, there is much left to study within channel-based mechanisms. For example, one open question is what is the most number of outcomes that can be semi-shattered by an agent with c channels? Another question in this domain is whether or not strict improvements in our upper bound on expected efficiency can be guaranteed as channels are added to the best channel-based mechanism.

We also believe that the efficiency bound and expressiveness measures we considered can be used to provide a richer view of the flaws of inexpressive mechanisms in a wide variety of domains, as we will begin to show in the next three chapters. For example, given the prior over the agents' types we can compute (or approximate) the likely loss in efficiency that will result from mechanisms with varying levels of expressiveness (of course this only provides a bound on this loss for a given mechanism, to compute the loss exactly we would need to extend our analysis to consider the mechanism's actual equilibrium).

In another direction, we can develop algorithms that take as input the prior over the agents' types in the particular setting at hand and output the efficiency-maximizing mechanism subject to a limit on expressiveness. (There has been significant work on developing algorithms for automated mechanism design in other settings [45–50, 68, 92, 124, 127–129, 131].)

This objective can be pushed even further to develop a methodology for identifying ways in which existing inexpressive mechanisms can be made more expressive to garner the greatest efficiency increase. For example, it may be possible to develop an algorithm which takes as input the prior over agent types, the maximum allowed expressiveness, and a default mechanism. This algorithm could then provide suggestions about how the default mechanism should be raised to the desired expressiveness level in a way that provides the largest improvement in its expected efficiency.¹¹

Finally, it has often been observed in practice that increases in expressiveness lead to increases in user burden because the increase in expressiveness is typically associated with an increase in the number and/or complexity of “queries” the user has to answer. However, more expressive mechanisms typically eliminate much (or all) of the strategic complexity (e.g., the cognitive effort required to speculate and counter-speculate about the strategies

¹¹One way to operationalize this idea is to search for a payment rule that forces agents as close as possible to implementing our upper bound on efficiency. A related problem involves assigning a limited number of channels to agents in a channel-based mechanism to optimize efficiency.

of other agents) that arises when agents are forced to shoehorn their preferences into an inexpressive mechanism. It may be possible to extend our theoretical framework to capture this tradeoff and explore the relationship between these two types of complexity in a variety of settings.

2.8 Appendix: Proofs of all technical claims

Proof of Proposition 1. Given a mechanism with reportable type space in \mathfrak{R}^d , we can construct an equivalent mechanism with reportable type space in \mathfrak{R} with an injective mapping from \mathfrak{R}^d to \mathfrak{R} . Then, when an agent makes a report in \mathfrak{R} , we use the reverse mapping and act as if the agent had expressed the corresponding point in \mathfrak{R}^n to the original mechanism.

One way to construct the injective mapping is as follows. Let σ_i^j be the i th bit (or digit) of the real number that the agent expresses for dimension $j \in \{1, 2, \dots, n\}$. Let p_k be the k th prime number. Our desired number in \mathfrak{R} is given by,

$$\prod_i \prod_j (p_{(i-1)n+j})^{\sigma_i^j}.$$

□

Proof of Proposition 2. When designing an outcome function, every time we increase a limit on the number of outcomes that agent i can (semi-)shatter, we can construct the function so that the agent can distinguish among all of the impact vectors it had previously distinguished between, plus at least one additional impact vector (the impact vector that was preventing it from (semi-)shattering the additional outcome). □

Proof of Proposition 3. The number of possible impact vectors for agent i with k different outcomes when the other agents have types T_{-i} is $|T_{-i}|^k$. Being able to shatter the full outcome set of size k requires that the agent be able to distinguish among each of these vectors, thus its maximum impact dimension must be greater than or equal to this amount. □

Proof of Proposition 4. The following reasoning demonstrates that Equation 2.2 is a valid upper bound on the maximum attainable expected efficiency by any mechanism using the outcome function f .

$$\begin{aligned}
E_{t^n} [\mathcal{E}(f)]^+ &= \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(m(t^n) = \theta^n) W(f(\theta^n), t^n) \\
&\leq \max_{B(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) \int_{\theta^n \in \Theta^n} P(B(t^n) = \theta^n) W(f(\theta^n), t^n) \\
&= \max_{\hat{B}(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) W(f(\hat{B}(t^n)), t^n) \\
&= \max_{\hat{B}(\cdot)} \int_{t^n \in T^n} P(T^n = t^n) W(f(\{\hat{B}_1(t_1), \dots, \hat{B}_n(t_n)\}), t^n).
\end{aligned}$$

The step between the second and third equations follows from the fact that one of the maxima of the function in the second equation must have each entry of $B(\cdot)$ (a function that maps every type vector to a mixed strategy profile) as a point mass. There is at least one single pure strategy combination for each type vector that leads to the outcome with highest welfare, thus there is no reason to consider mixed strategies in this bound. The last step is valid because the strategy of each agent can depend only on its own private type. \square

Proof of Proposition 5. First we will prove the forward implication, namely that the upper bound reaches full efficiency if any agent can distinguish among each of the impact vectors in at least one of its fully efficient sets.

The fact that some agent, i , can distinguish among each of the impact vectors in some fully efficient set, G_i^* , implies that there is a pure strategy for agent i , h_i , which is a mapping from its types to expressions, and a pure strategy profile for the agents other than i , h_{-i} , mapping from each of their types to expressions, that causes the most efficient outcome to be chosen by the mechanism for every possible combination of types. If we set $\hat{B}(t^n) = \{h_i(t_i), h_{-i}(t_{-i})\}$, then $E[\mathcal{E}(f)]^+$ will reach full efficiency.

Now we will prove the backward implication, namely that if any agent cannot distinguish among each of the impact vectors in at least one of its fully efficient sets, then the upper bound cannot reach full efficiency.

Let agent i be an agent that cannot distinguish among each of its impact vectors in any of its fully efficient sets. Consider any set of impact vectors that agent i can distinguish among, G_i . Based on the premise of the proposition, at least one of the impact vectors, g_i^* , corresponding to the fully efficient outcomes when agent i has type t_i^* , cannot be expressed

by agent i . This means that no matter which strategies the agents other than i choose, at least one of the outcomes chosen by the mechanism when agent i has type t_i^* will be less than fully efficient. \square

Proof of Proposition 6. We know that the premise implies there is some pure strategy for agent i , h_i , that achieves full efficiency when played against some pure strategy profile, h_{-i} , for the other agents. Let h_j be agent j 's pure strategy in the profile h_{-i} . Construct a new pure strategy profile, h_{-j} , by starting with h_{-i} and removing agent j 's pure strategy. Now add agent i 's pure strategy, h_i , to complete the profile. Since we have not changed the strategies played in any circumstances, h_j will achieve full efficiency against h_{-j} , thus completing our proof. \square

Proof of Proposition 7. In these settings, as soon as agent i knows its own type it knows for certain the single most efficient outcome. It never needs to distinguish among more than one-dimensional impact vectors and there are only $|\mathcal{O}|$ such vectors. \square

Proof of Corollary 1. This follows directly from Propositions 5 and 7. \square

Proof of Theorem 1. The set of mechanisms allowing agent i maximum impact dimension d_i is a super-set of the mechanisms allowing agent i maximum impact dimension $d_i' < d_i$. Thus, the fact that the bound for the best mechanism increases weakly monotonically is trivially true for any increase in d_i . The challenge is proving the strictness of the monotonicity.

Consider increasing d_i from $d_i^{(1)} < d_i^*$ to $d_i^{(2)} > d_i^{(1)}$. Let $G_i^{(1)}$ be the best (for efficiency) set of impact vectors that agent i can distinguish among when restricted to $d_i^{(1)}$ vectors (i.e., the set of $d_i^{(1)}$ impact vectors that maximize the upper bound on expected efficiency). We know that there are at least $d_i^* - d_i^{(1)} \geq 1$ impact vectors corresponding to fully efficient sets of outcomes that cannot be expressed by agent i , and thus at least that many fully-efficient impact vectors are absent from $G_i^{(1)}$. When we increase our expressiveness limit from $d_i^{(1)}$ to $d_i^{(2)}$, we can add one of those missing vectors to $G_i^{(1)}$ to get $G_i^{(2)}$. Since $G_i^{(2)}$ allows agent i to distinguish among all the same vectors as $G_i^{(1)}$ and an additional vector which corresponds to a fully efficient set of outcomes, the new mechanism with maximum impact dimension $d_i^{(2)}$ has a strictly higher expected efficiency bound. \square

Proof of Corollary 2. This follows directly from Theorem 1 and Proposition 2. \square

Proof of Lemma 1. Start with any number of outcomes and any number of types for the agents other than i with equal likelihood (and let the probability of any particular set of types for the agents other than i be independent of i 's type). Choose a set, G_i , of unique impact vectors for agent i with size d_i . Construct one non-zero probability type for agent i for each impact vector in G_i , $t_i^{g_i}$. Set the total welfare of all agents, as shown below, to an arbitrarily large number for every combination of joint types corresponding to the impact vectors in G_i (in an interdependent values setting there are no restrictions on the agent's utility functions, so the welfare function for each set of joint types can be constructed arbitrarily):

$$\forall g_i \in G_i, \forall t_{-i}, \quad W(g_i(t_{-i}), \{t_i^{g_i}, t_{-i}\}) = M.$$

If agent i cannot distinguish among all of the d_i impact vectors, then the efficiency bound will be arbitrarily smaller than if it could. Thus, for the best outcome function, the move from $d_i - 1$ to d_i results in an arbitrary increase in the bound on efficiency. \square

Proof of Lemma 2. The part that applies to the interdependent values setting follows directly from Lemma 1, since decreasing k_i by one also decreases d_i by at least 1.

To prove the implication for private values, we will construct a setting (i.e., utilities, types, and outcomes), such that agent i must be able to semi-shatter an outcome space of size k_i in order to avoid the upper bound being arbitrarily lower than full efficiency. Our constructed setting can have any number of outcomes, any number of other agents, and any number of joint types for the other agents. However, in order to assign the total utility of the other agents for each of their joint types in an arbitrary way, we will limit every other agent except for one, agent j , to a single type (agent j will have $|T_{-i}|$ types). We will set the utility of every agent other than i and j to 0 in all circumstances and build our construction using only these two agents.

We will start with a set of outcomes, \mathcal{O}' , that has size k_i (if $k_i = 1$ the rest of this proof is trivial, if every single outcome provides an arbitrary amount of welfare then not being able to make any one of them happen will lead to arbitrary inefficiency). We will assume, without loss of generality, that the outcomes in \mathcal{O}' are the only outcomes that any of the agents derive utility from. We will also assume that there is some strict ordering on the outcomes, from o_1 to o_{k_i} , and on agent j 's types, from $t_j^{(1)}$ to $t_j^{(|T_j|)}$.

We will now describe how to set the utility of agent j for every outcome under every one of its types. Our construction sets agent j 's utility for any outcome, o_m , under each of its types to be arbitrarily larger than for the outcome preceding it in the strict ordering, o_{m-1} , with the first outcome always leading to utility 0. For a fixed one of agent j 's types, all of the differences in utility for successive outcomes will be the same size. However, the gap amount (i.e., the difference in utility between o_m and o_{m-1}) will increase by an arbitrary amount for each successive successive type. This results in agent j 's utility under each of its types being a step function over the strictly ordered outcomes in \mathcal{O}' , with the step sizes increasing for each successive type. Formally, we will set agent j 's utility function in the following way (let M be an arbitrarily large number),

$$(\forall m, \forall l) \quad v_j(o_l, t_j^{(m)}) = (l - 1) \times ((m - 1) \times 2 \times M).$$

Now for each of the $\binom{|\mathcal{O}'|}{2}$ unordered pairs of outcomes, o_a and o_b (where a is always before b in our strict ordering), we will construct a set of $|T_j|$ types for agent i , which we will call $T_i^{(a,b)}$. Agent i 's utility under all of the types in $T_i^{(a,b)}$ will be hugely negative for all outcomes other than o_a and o_b (this value does not have to be negative infinity, just arbitrarily lower than the total welfare of any outcome under any circumstance), thus causing an arbitrary loss of efficiency if either of these outcomes is not chosen. Again, we will assume a strict ordering on the types in $T_i^{(a,b)}$, from 1 to $|T_j|$. Agent i 's utility for o_b under each of the types in $T_i^{(a,b)}$ will be set to the arbitrarily large number M , and for o_a (the typically less preferred outcome by agent j since it comes earlier in the ordering) will be set to successively increasing multiples of the distance between the outcomes in the strict ordering times twice the arbitrarily large number used above, (i.e., $(b - a) \times 2 \times M$). In other words, o_a will provide successively more utility to agent i as its type increases from 1 to $|T_j|$. Formally, we will set agent i 's utility under the types in $T_i^{(a,b)}$ to be the following,

$$\begin{aligned} & \left(\forall m \mid t_i^{(m)} \in T_i^{(a,b)} \right) \quad v_i(o_b, t_i^{(m)}) = M \\ & \left(\forall m \mid t_i^{(m)} \in T_i^{(a,b)} \right) \quad v_i(o_a, t_i^{(m)}) = (m - 1) \times (b - a) \times 2 \times M \\ & \left(\forall o_j \in \mathcal{O} \setminus \mathcal{O}', \forall m \mid t_i^{(m)} \in T_i^{(a,b)} \right) \quad v_j(o_j, t_i^{(m)}) = -\infty. \end{aligned}$$

When $t_i^{(m)}$ is matched with $t_j^{(m)}$, the total welfare of outcome o_b will be at least M larger than the total welfare of o_a . However, for all of j 's types smaller than m the opposite will

be true.

$$\begin{aligned} W(o_b, \{t_i^{(m)}, t_j^{(m)}\}) &= M + [(b-1) \times (m-1) \times 2 \times M] \\ W(o_a, \{t_i^{(m)}, t_j^{(m)}\}) &= [(b-a) \times (m-1) \times 2 \times M] + [(a-1) \times (m-1) \times 2 \times M] \\ &= [(b-1) \times (m-1) \times 2 \times M]. \end{aligned}$$

By constructing the utility functions in this way we have guaranteed that for any pair of agent j 's types, $t_j^{(m)}$ and $t_j^{(m')}$ (where $m < m'$ in our strict ordering), there is a type for agent i that requires o_b to happen against $t_j^{(m)}$ and o_a against $t_j^{(m')}$ to avoid an arbitrary loss in efficiency (because any other outcome would lead to at least M less welfare).

Now, we can repeat this process for each pair of outcomes in \mathcal{O}' by constructing types for agent i that select the pair. This guarantees that agent i must be able to make every pair of outcomes happen against every pair of agent j 's types in the same order, or else face an arbitrary loss of efficiency in some non-zero probability combination of types. This is equivalent to saying that agent i must be able to semi-shatter the outcome space \mathcal{O}' in order to avoid an arbitrary decrease in the expected efficiency bound. \square

Proof of Lemma 3. Let agent i 's utility for outcomes o_1 and o_2 under type $t_i^{(1)}$ and $t_i^{(2)}$ be denoted as X and Y , respectively. For the agents other than i , let the sum of their utilities for the outcomes o_1 and o_2 under types $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$ be denoted as a and b , and a' and b' , respectively. We wish to show that the ordering on efficient outcomes imposed by this collection of types cannot be reversed. Formally,

$$\begin{aligned} (X + a > Y + b) \text{ and } (Y + b' > X + a') &\Rightarrow \\ \neg(\exists X', Y') (X' + a < Y' + b) \text{ and } (Y' + b' < X' + a'). & \end{aligned}$$

We will proceed by assuming this is true, namely that there exists an X' and Y' that satisfy the second set of inequalities, and show that it leads to a contradiction. If all of the inequalities implied by this assumption held we would have the following,

$$\begin{aligned} b - a < X - Y < b' - a' \\ b' - a' < X' - Y' < b - a, \end{aligned}$$

Contradiction. \square

Proof of Theorem 2. The forward implication in both settings follows directly from Lemma 2. The backward implication in the interdependent values setting follows from Lemma 1 and Proposition 5 (since there will always be a fully efficient set that contains every possible impact vector). In the private values setting, the backward implication is implied by Lemma 3, since it proves that it is never necessary for full efficiency in such settings to shatter any set of outcomes (only to semi-shatter them). \square

Proof of Theorem 3. Let h_I^* be a pure strategy profile that achieves the expected efficiency of the bound $E[\mathcal{E}(f)]^+$. For shorthand, let $h^*(t_i)$ be agent i 's expression under type t_i and the pure strategy profile h^* , and let $h^*(t_{-i})$ denote the expressions of the agents other than i under that profile.

Consider the class of payment functions, π^+ , that charges agent i some constant function of the other agent's expressions minus the expected welfare of the other agents, given that agent i expresses θ_i and the other agents play the pure strategies denoted by h^* ,

$$\pi_i^+(\theta_i, \theta_{-i}) = C_i(\theta_{-i}) - E_{t_{-i}}[W_{-i}(f(\theta_i, h^*(t_{-i})), t_{-i})].$$

Now, we will prove that under any payment function in the class π^+ the pure strategy profile h^* is a Bayes-Nash equilibrium. The following inequality implies that it is always (weakly) preferable in expectation over the types of the agents other than i for agent i , under any type t_i , to report $h^*(t_i)$ rather than a different expression, assuming that the other agents play according to h^* as well. (In the first equation, agent i 's payment is outside of the expectation since it does not depend on the types of the other agents, and we omit the C_i terms since they do not depend on agent i 's expression.)

$$E[v_i(f(h^*(t_i), h^*(t_{-i})), t_i)] - \pi_i^+(h^*(t_i), h^*(t_{-i})) \geq E[v_i(f(\theta'_i, h^*(t_{-i})), t_i)] - \pi_i^+(\theta'_i, h^*(t_{-i}))$$

$$\begin{aligned} E[v_i(f(h^*(t_i), h^*(t_{-i})), t_i)] + E[W_{-i}(f(h^*(t_i), h^*(t_{-i})), t_{-i})] &\geq \\ E[v_i(f(\theta'_i, h^*(t_{-i})), t_i)] + E[W_{-i}(f(\theta'_i, h^*(t_{-i})), t_{-i})]. \end{aligned}$$

The left-hand side of the final inequality is the expected welfare when the agents play the pure strategy profile h^* and agent i has type t_i . The right-hand side is the expected welfare when agent i deviates from h^* under type t_i . This inequality holds because it is impossible

for any deviation from h^* to increase expected welfare. Based on our assumptions, it would have to already be reflected in h^* . \square

Proof of Proposition 8. We can set the C_i term in π_i^+ for each agent to be the average expected total payment to all other agents. This amount does not depend on the agent's own expression, and it gives strong budget balance because each agent pays an equal share of the total payments made to the other agents.

The following reasoning proves that, assuming the preference distribution satisfies the non-negative externality criterion for the given outcome function, the resulting mechanism is *ex ante* individually rational (i.e., that any agent's utility for participating is always positive in expectation, prior to learning its own type). $E[u_i]$, the expected utility of agent i , is given by the following.

$$\begin{aligned} E_t[u_i] &= E_t [v_i(f(h^*(t_i), h^*(t_{-i})), t_i) - \pi_i^+(h^*(t_i), h^*(t_{-i}))] \\ &= E_t \left[W(f(h^*(t)), t) - \frac{1}{(n-1)} \left(\sum_{j \neq i} \sum_{k \neq j} v_k(f(h^*(t_k), h^*(t_{-k})), t_k) \right) \right] \\ 0 &\leq E_t \left[W(f(h^*(t)), t) - \left(v_i(f(h^*(t)), t_i) + \frac{(n-2)}{(n-1)} W_{-i}(f(h^*(t)), t_{-i}) \right) \right] \end{aligned}$$

\square

Proof of Theorem 4. We will prove this by showing that the outcome function implementing our bound under a limit on expressiveness does not necessarily satisfy the weak-monotonicity (W-Mon) property, which has been shown to be a necessary condition for dominant-strategy implementation [25].

Consider the following example where agent one has three types and agent two has two types. The agents' valuations for each of three different outcomes are given below.

{	Agent 1	type	A	B	C
	$t_1^{(1)}$	14	0	0	
	$t_1^{(2)}$	0	0	0	
	$t_1^{(3)}$	1	0	12	

{	Agent 2	type	A	B	C
	$t_2^{(1)}$	11	13	0	
	$t_2^{(2)}$	10	0	0	

Under the valuations given above, the total social welfare of each outcome is given by the following table (the welfare of the most efficient outcome associated with each joint type is shown in bold).

Outcome	$t_1^{(1)}, t_2^{(1)}$	$t_1^{(1)}, t_2^{(2)}$	$t_1^{(2)}, t_2^{(1)}$	$t_1^{(2)}, t_2^{(2)}$	$t_1^{(3)}, t_2^{(1)}$	$t_1^{(3)}, t_2^{(2)}$
A	25	24	11	10	12	11
B	13	0	13	0	13	0
C	0	0	0	0	12	12

Consider a direct-revelation mechanism with a socially optimal outcome function, f . The impact vectors with the highest social welfare for agent one correspond to $[A, A]$, $[B, A]$, and $[B, C]$ (these are the outcomes with the greatest welfare under each combination of types). If we are forced to design an outcome function that limits agent one to maximum impact dimension $d_1 \leq 2$ and the type $t_1^{(2)}$ is highly unlikely (e.g., $P(t_1^{(2)}) = \epsilon$), then the outcome function with the highest expected welfare will provide agent one with the impact vectors, $[A, A]$, $[A, A]$ and $[B, C]$.

The W-Mon property states that the following inequality must hold for all t_1, t'_1 , and t_2 ,

$$v_1(f(t_1, t_2), t_1) - v_1(f(t'_1, t_2), t_1) \geq v_1(f(t_1, t_2), t'_1) - v_1(f(t'_1, t_2), t'_1).$$

If we use $t_1^{(2)}$ and $t_1^{(3)}$ for t_1 and t'_1 , respectively, and $t_2^{(1)}$ for t_2 , then we can rewrite the inequality for our limited-expressiveness mechanism as follows,

$$\begin{aligned} v_1(f(t_1^{(2)}, t_2^{(1)}), t_1^{(2)}) - v_1(f(t_1^{(3)}, t_2^{(1)}), t_1^{(2)}) &\geq v_1(f(t_1^{(2)}, t_2^{(1)}), t_1^{(3)}) - v_1(f(t_1^{(3)}, t_2^{(1)}), t_1^{(3)}) \\ v_1(A, t_1^{(2)}) - v_1(B, t_1^{(2)}) &\geq v_1(A, t_1^{(3)}) - v_1(B, t_1^{(3)}). \end{aligned}$$

This inequality is violated by the valuation functions in our example, so the inexpressive mechanism cannot be implemented in dominant strategies. \square

Proof of Proposition 9. Suppose the opposite: that agent i has maximum impact dimension d_i under f , and there exists an outcome function, f' , that emulates f while providing agent i with less than d_i expressions. Let G_i be one of the largest sets of impact vectors that the agent can distinguish among under f (i.e., $|G_i| = d_i$ and $D_i(G_i)$ is true), let h_{-i} be a strategy by the other agents that allows i to distinguish among the vectors in G_i under f , and let q be the mapping that allows f' to emulate f .

Since d_i expressions are needed to distinguish among each of the different impact vectors in G_i , our assumption implies that there will be at least two distinct impact vectors in G_i that agent i cannot distinguish among under f' . Let these be denoted as $g_i^{(1)}$ and $g_i^{(2)}$. Since $g_i^{(1)}$ and $g_i^{(2)}$ are distinct, there must be at least one joint type for the agents other than i , t_{-i} , such that they map to different outcomes. Furthermore, the impact vectors are expressible under f , so it must be possible for agent i to cause both the outcome mapped by $g_i^{(1)}$ and the outcome mapped by $g_i^{(2)}$ under t_{-i} to be chosen by f (i.e., there exists a $\theta_i^{(1)}$ and $\theta_i^{(2)}$ such that $f(\theta_i^{(1)}, h_{-i}(t_{-i})) \neq f(\theta_i^{(2)}, h_{-i}(t_{-i}))$).

However, since agent i cannot distinguish between $g_i^{(1)}$ and $g_i^{(2)}$ under f' , $\theta_i^{(1)}$ and $\theta_i^{(2)}$ must map to the same expression under q . Thus, we get the following starting from the equation above,

$$\begin{aligned} f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))) &\neq f'(q_i(\theta_i^{(2)}), q_{-i}(h_{-i}(t_{-i}))) \\ f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))) &\neq f'(q_i(\theta_i^{(1)}), q_{-i}(h_{-i}(t_{-i}))). \end{aligned}$$

Contradiction. □

Proof of Lemma 4. If one of the agents other than i lies, the number of impact vectors that agent i can distinguish among can only decrease. Any impact vector that was distinct because of an outcome chosen under the type that is now being reported untruthfully will no longer be distinct, and no new impact vectors can become distinct because of the lie. □

Proof of Proposition 10. We have already shown in Proposition 9 that no outcome function can emulate f using fewer expressions for each agent, i , than its maximum impact dimension under f , d_i . We will now show that if f is a direct-revelation outcome function, it is always possible to emulate it using exactly d_i expressions for each agent. To prove this, we will

use Lemma 4, which implies that we can construct a mapping from expressions under f to expressions under f' such that any two types resulting in the same impact vector under f are also mapped to the same expression under f' . We can set the outcomes in f' as follows,

$$f'(q_i(t_i), q_{-i}(t_{-i})) = f(t_i, t_{-i}).$$

This will produce a valid mapping because the types that map to to the same expression will result in the same outcome for all joint types of the other agents. \square

Proof of Corollary 3. From Theorem 1, we know that this is true for maximum impact dimension. From Proposition 9, we know that the set of outcome functions that limit agent i to d_i expressions does not include any outcome functions where agent i has maximum impact dimension greater than d_i . A direct-revelation mechanism can always be designed to maximize the bound subject to a limit on expressiveness, and that outcome function can be emulated by one that provides d_i expressions to each agent i . \square

Proof of Corollary 4. This follows directly from Theorem 2, which states that an agent may need to shatter its entire outcome space to avoid arbitrary inefficiency, and Proposition 3, which states that $|T_{-i}|^{|\mathcal{O}|}$ expressions are needed needed to shatter an outcome space. \square

Proof of Lemma 5. Let Θ_{-i} be expressions for the agents other than i that allow agent i to shatter T_{-i} (if f is a direct-revelation mechanism these two sets will be identical). Construct a total ordering of Θ_{-i} so that for any pair, θ_{-i}^j and θ_{-i}^k (where $j < k$), and any expression by agent i , $\theta_i \in \Theta_i$, that causes the mechanism to choose A and B when the other agents express θ_{-i}^j and θ_{-i}^k , A is chosen for θ_{-i}^j and B for θ_{-i}^k . If this condition is not met, we can simply switch j and k . Re-labeling all of the expressions to satisfy this condition is possible because of the semi-shattering requirement that, under any strict ordering of the expressions of the agents other than i , all expressions by agent i in Θ_i that cause outcomes A and B to be chosen by f do so in the same order.

Consider a subset of expressions by agent i , Θ'_i , that allow it to choose between A and B for each of the immediately subsequent, or neighboring, pairs of expressions by the agents other than i . In other words, there exists at least one $\theta'_i \in \Theta'_i$, such that f chooses A and B under the expression pairs $(\theta'_i, \theta_{-i}^j)$ and $(\theta'_i, \theta_{-i}^{j+1})$, for all j . Additionally, order Θ'_i so that an expression that chooses between A and B when the other agents express θ_{-i}^1 and θ_{-i}^2 is the

first in the ordering. This expression is then followed in the ordering by one that chooses between A and B when the other agents express θ_{-i}^2 and θ_{-i}^3 , and so on until the last of the neighboring pairs of θ_{-i} 's is reached.

Under this ordering, none of the $|T_{-i}| - 1$ pairs of expressions along the diagonal of an input matrix corresponding to agent i 's expressions in Θ'_i (i.e., pairs of the form $(\theta_i^j, \theta_{-i}^j)$) can be in the same A -monochromatic rectangle. To see this, consider that whenever θ_i^j is matched against an expression larger than θ_{-i}^j , the outcome function does not choose A (since we ensured that A 's always come before B 's). Thus, all of the A 's in the reduced input matrix corresponding to Θ_{-i} and Θ'_i are to the left of the diagonal (in a triangle pattern) and the diagonal is all A 's. This implies that each of the pairs to the left of the diagonal must be in a different A -monochromatic rectangle, since they all result in the same value but are different when crossed with any other member of the set. We can reverse the total ordering of Θ_{-i} and make the same argument for B -monochromatic rectangles. \square

Proof of Theorem 5. The upper bound follows directly from Propositions 9 and 10, since they imply that an agent's maximum impact dimension is an upper bound on the number of messages needed by the best communication protocol for running f . The lower bound follows from Lemma 5, since it implies that if agent i can semi-shatter a set of types containing some outcome o , T_{-i}^o , there must be at least $|T_{-i}^o| - 1$ monochromatic rectangles for outcome o in any partitioning of f 's input matrix from i 's perspective. It has been previously shown that each rectangle requires at least one message in any communication protocol [87]. \square

Proof of Proposition 11. We will prove this statement for the semi-shatterable outcome dimension, k_i , which will imply it is true for maximum impact dimension as well (based on Proposition 2).

Consider any channel-based mechanism that assigns c_i channels to agent i and allows it a semi-shatterable outcome dimension $k_i < |\mathcal{O}|$. We will assume from here on that $k_i \geq 2$, since if $k_i = 1$ the theorem is trivially true (we can build a fully expressive VCG mechanism over 2 outcomes with a single channel and thus adding a channel will definitely increase k_i to at least 2).

Let the largest set of outcomes that agent i can semi-shatter over in this mechanism be \mathcal{O}' . There is a non-empty set of outcomes missing from \mathcal{O}' , we will call that $\mathcal{O}^* = \mathcal{O} \setminus \mathcal{O}'$.

Now consider adding one channel for agent i to the mechanism and connecting it to one of the outcomes $o^* \in \mathcal{O}^*$. The agent can still semi-shatter over \mathcal{O}' , since it can just ignore the new channel. However, it can now also semi-shatter a larger set, $\mathcal{O}' \cup \{o^*\}$.

With the additional channel connected to o^* the agent can control the amount of utility it reports on this outcome arbitrarily (without affecting its reports on any other outcomes). Consider any pair of outcomes in the original set, $o'_1, o'_2 \in \mathcal{O}'$. Agent i can now make o^* happen against any type where either of those outcomes happened in the old mechanism by setting its report on the new channel to be ϵ greater than the sum of its reports on the channels connected to the outcome it used to select. Formally, if C_i is the channel mapping from the original mechanism, then we can translate any report in the old mechanism, θ_i , to a report in the new mechanism, θ_i^* , which causes o^* to happen whenever any other outcome, o' , did previously.

$$\begin{aligned} (\forall j \mid 1 \leq j \leq c_i) \theta_{i,j}^* &= \theta_{i,j} \\ \theta_{i,c_1+1}^* &= \sum_{j \in C_i(o')} \theta_{i,j} + \epsilon. \end{aligned}$$

Since agent i can do this with both outcomes from the original semi-shatterable set, we have confirmed that it has reports in the new mechanism that make o^* happen with every pair of outcomes in \mathcal{O}' (this is an inductive argument, since each of those outcomes had this property before).¹² Thus, agent i can semi-shatter the new larger outcome set using the additional channel. \square

Proof of Corollary 5. The fact that the bound is weakly monotonic is true because the extra channel can always be ignored. The fact that the increase can be arbitrarily large follows directly from Proposition 11 and Lemma 2 (since increasing the number of channels by one can be used to increase the agent's semi-shatterable outcome dimension). \square

¹²We have assumed the agent was not using the tie-breaking properties of the original mechanism to shatter the outcomes. If this assumption does not hold, the proof is still valid as long as the mechanism always breaks ties consistently (i.e., when the channels connected to outcomes o_1 and o_2 have the same sum it always chooses either o_1 or o_2).

Proof of Proposition 12. This proof is based on a pigeon hole argument. With fewer than $\lceil \log_2(|\mathcal{O}|) \rceil$ channels there will be at least two outcomes connected to the exact same set of channels. If agent i has C_i channels, then it has 2^{C_i} sets of channels. When C_i is small the number of sets of channels will be less than the number of outcomes,

$$C_i < \lceil \log_2(|\mathcal{O}|) \rceil \Rightarrow 2^{C_i} < |\mathcal{O}|.$$

This will prevent the agent from forcing the mechanism to choose both of those outcomes with different expressions, since the agent's own contribution to the two outcomes will always be identical. \square

Proof of Proposition 13. We will show that no agent can shatter any set of two outcomes against any two types, even when it has a channel dedicated solely to each of the two outcomes (so that it can place an arbitrary amount of value on either outcome). This implies that it is impossible to shatter any larger set of outcomes or types in any channel-based mechanism.

We will assume, for contradiction, that there is some agent i that can shatter a pair of outcomes A and B in a channel-based mechanism. Let agent i 's total channel value connected to outcome A be X and let its total channel value connected to B be Y . Consider two types for the agents other than i , $t_{-i}^{(1)}$ and $t_{-i}^{(2)}$, and the reports mapped to them under *any* pure strategy, $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$. Let the sum of the reports by the other agents on the channels connected to A be denoted a_1 and a_2 under the first and second expressions, respectively. Likewise, let b_1 and b_2 be the sum of the reports on B . We have assumed (for contradiction) that there exists an X, Y, X' and Y' that satisfy the following inequalities.

$$\begin{aligned} A \text{ against } 1, B \text{ against } 2 & \begin{cases} X + a_1 > Y + b_1 \\ Y + a_2 > X + b_2 \end{cases} \\ B \text{ against } 1, A \text{ against } 2 & \begin{cases} Y' + b_1 > X' + a_1 \\ X' + a_2 > Y' + b_2. \end{cases} \end{aligned}$$

This leads directly to the following.

$$\begin{aligned} b_1 - a_1 &< X - Y < b_2 - a_2 \\ b_2 - a_2 &< X' - Y' < b_1 - a_1. \end{aligned}$$

Contradiction. □

Proof of Corollary 6. This follows directly from Proposition 13 and Lemmas 1 and 2. □

Proof of Proposition 14. With $|\mathcal{O}| - 1$ channels, we can construct a VCG outcome function in the following manner. For each agent i , connect each of i 's channels to a different outcome, leaving one outcome with no channel from that agent. The agent then reports its utility under each outcome relative to the outcome with no channels. The mechanism chooses outcome whose channels have the largest sum, which is equivalent to choosing the welfare-maximizing outcome. The payment rule will not be affected by the fact that each agent is reporting its utility relative to a particular outcome. To see this, consider the VCG (i.e., Clarke tax) payment of any agent i . This payment is equal to the total difference in utility of the other agents, had agent i not participated. Let the outcome with agent i in the mechanism be A , and the outcome without agent i be B . Let the outcome with no channels attached be o_j , for every agent j . Then we have the payment for agent i as,

$$\begin{aligned} \pi_i &= \sum_j (v_j(A, t_j) - v_j(o_j, t_j)) - \sum_j (v_j(B, t_j) - v_j(o_j, t_j)) \\ &= \sum_j (v_j(A, t_j) - v_j(B, t_j)) - \sum_j (v_j(o_j, t_j) - v_j(o_j, t_j)) \\ &= \sum_j (v_j(A, t_j) - v_j(B, t_j)). \end{aligned}$$

Since the $v_j(o_j, t_j)$ terms drop out of this equation, having every agent report their utility for every outcome minus their utility for one particular outcome does not effect the calculation. This shows that the payment rule can be properly calculated even when each agent has a single outcome with no channels.

Using a pigeon hole argument, we can see that an agent with fewer than $|\mathcal{O}| - 1$ channels will either have at least two sharing a channel, making it impossible for that agent to express arbitrary non-linear utility for every outcome (a requirement for implementing a VCG mechanism), or it will have two outcomes without a channel, making it impossible for that agent to express any preference for one of the outcomes. □

Proof of Lemma 6. We will prove only the forward implication. Once that is proved, the backward implication will be trivial since we can just switch the labels of C and D . From

the premise, we have $(A \setminus C = D \setminus B)$, which implies that every element, x , in A and not in D is in either B or C , and every element, y , in A and B must also be in C (if y were not in either C or D , or if y were in only D , it would contradict the premise). Thus, the difference between A and D must be contained completely in C (i.e., $(A \setminus D) \subseteq C$). The following reasoning proves the rest of the claim,

$$\begin{aligned} A \setminus D &\subseteq C \\ A \setminus D &= C \setminus (C \setminus A) \\ &= C \setminus (B \setminus D) \\ &= C \setminus B. \end{aligned}$$

The last step is valid because we know that no elements from D can be in the set on the right-hand side, once all the operators are applied (since the left-hand side involves removing all elements in D from A). Thus, it cannot make a difference if we leave them in B before subtracting it from C , since the set minus operator in the parentheses on the right-hand side only serves to maintain the elements from D in the resulting set. This same logic can be repeated for the other half of the conjunction in the premise. \square

Proof of Lemma 7. From Lemma 6, in addition to our premise, we know that the following must also be true (we drop the i subscript on the channel sets for shorthand, since all sets of channels discussed in this proof belong to agent i),

$$(S^A \setminus S^D = S^C \setminus S^B) \quad \text{and} \quad (S^D \setminus S^A = S^B \setminus S^C).$$

We will assume, for contradiction, that agent i can semi-shatter both pairs of outcomes, $\{A, B\}$ and $\{C, D\}$. From Observation 1, we know that in order for i to be able to semi-shatter a set of outcomes, it must be able to semi-shatter it for any *pair* of types of the other agents. Thus, there must be at least one pair of reports by the agents other than i , $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, such that agent i can cause all four outcomes to happen (although we are considering semi-shattering, so the order in which they happen does not matter). Let the sum of the reported channels under the first (second) profile for the other agents connected to outcome A be a_1 (a_2), to outcome B be b_1 (b_2), and so on.

We will assume, without loss of generality, that $b_1 - a_1 < b_2 - a_2$ and that A will happen against $\theta_{-i}^{(1)}$ and B will happen against $\theta_{-i}^{(2)}$ (if the inequality does not hold, we can reverse

the labels on the θ_{-i} 's). In order to cause A to happen against the first opponent profile, and B against the second, the following inequalities must hold (from here on we use the shorthand S^A to denote the sum of agent i 's report on the channels in S^A , and we assume that ties are broken consistently so that an agent cannot use them to semi-shatter).

$$A \text{ happens against } 1 \begin{cases} S^A + a_1 > S^B + b_1 \\ S^A + a_1 > S^C + c_1 \\ S^A + a_1 > S^D + d_1 \end{cases}$$

$$B \text{ happens against } 2 \begin{cases} S^B + b_2 > S^A + a_2 \\ S^B + b_2 > S^C + c_2 \\ S^B + b_2 > S^D + d_2 \end{cases}$$

Let the difference between the sum of channels in S^A and S^C be denoted S_1 (i.e., $S^A - S^C = S_1$). From the premise, we have that $S^D - S^B = S_1$. This is because the channels that are in S^A and not S^C are the same as those that are in S^D and not S^B . Additionally, the channels in S^C that are not in S^A are the same as those that are in S^B and not S^D . Let the difference in the sum of the channels in S^A and S^D be denoted by S_2 . This leads the following equality, which is implied by Lemma 6: $S^A - S^D = S^C - S^B = S_2$. Now the equations above simplify to the following.

$$\begin{aligned} b_1 - a_1 &< S^A - S^B < b_2 - a_2 \\ c_1 - a_1 &< S_1 < b_2 - d_2 \\ a_1 - d_1 &< S_2 < b_2 - c_2 \end{aligned}$$

In order to semi-shatter C and D , with C happening against the first report by the other agents and D against the second, we have the following inequalities generated in the same fashion.

$$\begin{aligned} c_1 - d_1 &< S^C - S^D < c_2 - d_2 \\ b_2 - d_2 &< S_1 < c_1 - a_1 \\ b_1 - c_1 &< S_2 < a_2 - d_2 \end{aligned}$$

In order to semi-shatter over C and D in the opposite direction (with D first and C second) the constraints would change to the following.

$$\begin{aligned} c_2 - d_2 &< S^C - S^D < c_1 - d_1 \\ b_1 - d_1 &< S_1 < c_2 - a_2 \\ b_2 - c_2 &< S_2 < a_1 - d_1 \end{aligned}$$

Our assumption that agent i could semi-shatter both sets of outcomes under a single pair of types leads to a contradiction, since the following sets of constraints would have to be satisfied.

$$\begin{aligned} c_1 - a_1 &< b_2 - d_2 \\ b_2 - d_2 &< c_1 - a_1 \\ &\text{or,} \\ c_2 - b_2 &< a_1 - d_1 \\ a_1 - d_1 &< c_2 - b_2 \end{aligned}$$

Contradiction. □

Proof of Theorem 6. We will prove this by providing a distribution over valuations such that a channel-based mechanism that treats agent one's bid on any bundle Q to be the sum of its bids on some two other non-overlapping bundles, q_1 and q_2 , cannot achieve within 5% of the maximum expected efficiency.

We will first show that, in such a mechanism, agent one cannot choose between the pairs of outcomes where it wins q_1 or q_2 , and Q or nothing, since the channels connected to these outcomes overlap in the fashion described in Lemma 7. Let A be an outcome under which agent one is allocated bundle q_1 , let B be an outcome under which it is allocated q_2 , C for Q and D for nothing (also let S^A , S^B , S^C , and S^D be the sets of channels connected to those outcomes for the agent). Since agent one's bid on Q equals the sum of its bid on q_1 and q_2 , we have that $S^C = S^A \cup S^B$, and its bid for the outcome where it wins nothing is always 0, so we have $S^D = \emptyset$. These sets of channels meet the conditions of Lemma 7.

$$\begin{aligned} (S^A \setminus S^C = S^B \setminus S^D) &\quad \text{and} \quad (S^C \setminus S^A = S^D \setminus S^B) \\ (S^A \setminus (S^A \cup S^B) = \emptyset \setminus S^B) &\quad \text{and} \quad ((S^A \cup S^B) \setminus S^A = S^B \setminus \emptyset) \end{aligned}$$

Now, consider the following example with two agents, where each agent has two equally likely types: a “substitutable” type and a “complementary” type. The agents’ valuations for the bundles q_1 , q_2 , and Q are given below. Valuations for all other bundles, including the empty bundle, are assumed to be 0 or their minimum possible value. Since the items other than those in q_1 , q_2 , and Q provide no utility to either agent, and the bids for these items cannot affect how the items in q_1 , q_2 , and Q are allocated, we ignore the additional items in the rest of our proof.

Agent 1	type	q_1	q_2	Q
	t_1^s	0.5	0.5	0.5
	t_1^c	0	0	0.75

Agent 2	type	q_1	q_2	Q
	t_2^s	0	0.5	0.5
	t_2^c	0.75	0	1

Under the valuations given above, the total social welfare of each outcome is given by the following table (the welfare of the most efficient outcome associated with each joint type is shown in bold).

Outcome	t_1^s, t_2^s	t_1^s, t_2^c	t_1^c, t_2^s	t_1^c, t_2^c
$A : \{q_1, q_2\}$	1	0.5	0.5	0
$B : \{q_2, q_1\}$	0.5	1.25	0	0.75
$C : \{Q, \emptyset\}$	0.5	0.5	0.75	0.75
$D : \{\emptyset, Q\}$	0.5	1	0.5	1

The maximum expected efficiency, $E[\mathcal{E}^*]$, is then given by the following. We drop the t_1 and t_2 notation in favor of shorthand where types are simply referred to as s or c . P_{sc} denotes the probability of agent one having type s and agent two having type c .

$$\begin{aligned}
 E[\mathcal{E}^*] &= P_{ss}W(A, \{s, s\}) + P_{sc}W(B, \{s, c\}) + P_{cs}W(C, \{c, s\}) + P_{cc}W(D, \{c, c\}) \\
 &= \frac{1}{4} \times 1 + \frac{1}{4} \times 1.25 + \frac{1}{4} \times 0.75 + \frac{1}{4} \times 1 = 1.
 \end{aligned}$$

Since agent one cannot choose between the pairs of outcomes $\{A, B\}$ and $\{C, D\}$, the mechanism cannot achieve the expected efficiency of the optimal allocation for some combination

of types. In the best case, it will assign the second best outcome for one of the $\{s, c\}$, $\{c, s\}$, or $\{c, c\}$ types, which will cost at least 6.25% in expected efficiency. \square

Chapter 3

Expressiveness in Advertisement Auctions

3.1 Introduction

The sponsored search industry accounts for tens of billions of dollars in revenue annually. The most frequent variant of these auctions, the *generalized second price (GSP)* mechanism used by Google, Yahoo!, MSN, etc. solicits a single bid from each advertiser (i.e., agent) for a keyword and assigns the advertisers to positions on a search result page according to the bids (roughly speaking, with the first position going to the highest bidder, the second position to the second highest, etc.). Since agents cannot offer a separate bid price for each ad position, the GSP mechanism is fundamentally inexpressive, and more expressive variants have begun to receive more attention, (e.g., [33, 60, 89, 114]). In this chapter, we will attempt to characterize the loss of economic efficiency caused by this inexpressiveness, and to explore the conditions that affect that loss.

We begin by adapting our theoretical framework for studying expressiveness to analyze the GSP. We show that the notion of semi-shattering we introduced in Chapter 2 can capture the GSP's inexpressiveness, and we prove that for some preference distributions the GSP is arbitrarily inefficient.

However, in order to measure this inefficiency in practice we must be able to predict the outcome of the mechanism. The equilibrium of the GSP is known when it is assumed that agents have complete information (i.e., no private information about valuations) and monotonic preferences over positions (i.e., higher positions are always preferred) [149]; however when we relax these somewhat restrictive assumptions, the equilibrium behavior is unknown. In fact, it is often difficult to characterize equilibrium behavior in less than fully expressive mechanisms when agents have complex preferences [119, 139, 155]. For that reason, we develop a general tree search technique for computing an upper bound on a mechanism's expected efficiency that involves finding social welfare maximizing strategies for the agents. In the worst case, our search algorithm takes time that is exponential in the number of agents and types, but it can be applied to any preference distribution and provides an upper bound that tightens in an anytime manner.

We conclude with a series of experiments comparing the GSP to our slightly more expressive mechanism, which solicits an extra bid for premium ad positions, which we coin *Premium GSP (PGSP)*. We generate a range of realistic synthetic preference distributions based on published industry knowledge, and apply our search technique to compare the

efficiency bounds achieved by social welfare maximizing strategies in the two mechanisms. We also examine the performance of the two mechanisms when agents use a straightforward heuristic bidding strategy.

While we must be careful not to read too much into experiments on synthetic data, they suggest that the GSP’s efficiency loss can be dramatic. It is greatest in the practical case where some agents (“brand advertisers”) prefer top positions while others (“value advertisers”) prefer middle positions (since customers who click on ads in middle positions are more likely to take action, resulting in revenue). The loss is also worst when agents have small profit margins. Despite the fact that our PGSP mechanism is only slightly more expressive (and thus not much more cumbersome), it removes almost all of the efficiency loss in the settings we study empirically.

3.2 Setting and background results

The setting we study in this chapter (like most prior work, e.g., [59, 149]) is a one-shot auction for a set of k advertising positions that are ranked from 1 to k (rank 1 is the highest rank). In the model there are n agents. Each agent i has some private information (not known by the mechanism or any other agent) denoted by a type, t_i , (e.g., a vector of valuations, one for each of the k positions) from the space of the agent’s possible types, T_i .

Settings where each agent has a utility function, $u_i(t_i, O)$, that depends only on its own type and the outcome (matching of agents to positions), $O \in \mathcal{O}$, chosen by the mechanism are called *private values* settings. We also discuss more general *interdependent values* settings, where $u_i = u_i(t^n, O)$, i.e., an agent’s utility depends on the others’ private signals as well (for example, if one agent’s value for a position depends on market estimates of the other agents). In both settings, agents report expressions to the mechanism, denoted θ_i , based only on their own types. In the GSP mechanism each agent can report a single real value indicating his/her bid. A mapping from types to expressions is called a *pure strategy*.

Based on these expressions the mechanism computes the value of an outcome function, $f(\theta^n)$, which chooses an outcome. In the GSP mechanism, the outcome function maps agents to positions based on the order of their bids (the highest bidder is assigned the first position,

the second highest bidder is assigned the second, etc.).¹ The mechanism may also compute the value of a payment function, $\pi(\theta^n)$, which determines how much each agent must pay or get paid. In this chapter, we ignore the mechanism’s payment function because our notions of expressiveness are tied directly to a mechanism’s outcome function.²

As in Chapter 2, we denote by $W(t^n, o)$ the social welfare of outcome o when agents have private types (or private signals) t^n , i.e., $W(t^n, o) = \sum_i u_i(t^n, o)$.

3.2.1 A summary of our expressiveness theory framework

The theoretical framework that we developed in Chapter 2 provides the foundations for understanding the impact of making mechanisms more or less expressive, by providing meaningful, general definitions of a mechanism’s expressiveness.

In that chapter, we defined an *impact vector* to capture the impact of a particular expression by an agent under the different possible types of the other agents, and an expressiveness concept based on a notion called *shattering*, which we adapted from the field of computational learning theory [147]. The adapted notion captures an agent’s ability to distinguish among each of the impact vectors involving a subset of outcomes.

We also introduced a slightly weaker adaptation of shattering, called *semi-shattering*, for analyzing the more restricted setting where agents have private values. It captures an agent’s ability to cause each of the unordered pairs of outcomes (with replacement) to be chosen for every pair of types of the other agents, but without being able to control the *order* of the outcomes (i.e., which outcome happens for which type). In other words, there must exist a pair of fixed expressions made by the agents other than i such that agent i can cause any two outcomes to be chosen by varying its own expression. We defined a measure

¹In practice the bids are adjusted by predicted click-through rates (CTR) before conducting the ranking. For simplicity, we do not weight by CTR. However, our formulation can be easily extended to account for this by multiplying each agent’s original bid by its CTR.

²Since the efficiency bound that we study does not directly depend on equilibrium behavior, this is without loss of generality, as long as agents do not care about *each others’* payments. The equilibrium behavior of a given mechanism in practice may heavily depend on the payment function used. However, as we showed in Chapter 2, in all private values settings it is possible to design a payment function that implements this bound.

of expressiveness based on the size of the largest outcome space that an agent can shatter or semi-shatter. It is called the *(semi-)shatterable outcome dimension*.

In addition to defining the expressiveness notions, we tied those notions to an upper bound (Equation 2.2) on the expected efficiency of a mechanism’s *most efficient* equilibrium. We derived the bound by making the optimistic assumption that the agents play strategies which, taken together, attempt to maximize social welfare. Since we identify gaps in this bound due to reduced expressiveness with agents attempting to maximize social welfare, such a gap will exist under any strategy played by the agents. Chapter 2 provided several results relating this bound to a mechanism’s expressiveness. For the purposes of this chapter, Theorem 2 from Section 2.4, which proves that arbitrary inefficiency can result from the inability of an agent to semi-shatter a pair of outcomes, will prove useful.

3.3 Adapting our theory of expressiveness to ad auctions

In order to study the expressiveness properties of the GSP’s outcome function, we first derive a mathematical representation of the function. Let $R(i, o)$ be the rank of the position given to the i ’th agent in the matching of agents to positions denoted by outcome o . For analysis purposes, we will assume, without loss of generality, that each agent’s bid, θ_i , is restricted to be between 0 and 1 (this is not a limiting assumption due to the fact that we can losslessly map from any real valued space to this interval). Under this assumption, the following is functionally equivalent to the GSP’s outcome function.

$$(3.1) \quad f(\theta^n) = \arg \max_{o \in O} \sum_{i=1}^n (\theta_i \times 10^{-R(i,o)})$$

This function chooses the outcome that maximizes a weighted sum of the bids. Each bid in the sum is weighted by 10 raised to the negative power of the corresponding agent’s rank under the chosen outcome. Thus, agents with higher bids will contribute significantly more to the overall sum when they are placed in the first position, less when they are in the second, etc.³

³In fact, any weighting scheme can be used as long as lower ranking positions always have lower weights

We will now show that the outcome function of the GSP mechanism is inexpressive according to the notion of outcome semi-shattering introduced in the previous section.

Theorem 7. *Consider a set of outcomes, $\{A, B, C, D\}$, under which agent i is assigned different positions. In the GSP mechanism, agent i cannot semi-shatter both pairs of outcomes $\{A, B\}$ and $\{C, D\}$ if the other agents have more than one joint type and the ranks satisfy $R(i, A) < R(i, C) < R(i, D) < R(i, B)$.*

Proof. We will assume, for contradiction, that agent i can semi-shatter both pairs of outcomes, $\{A, B\}$ and $\{C, D\}$. Lemma 3 from Chapter 2 implies that there must be at least one pair of bids by the agents other than i , $\theta_{-i}^{(1)}$ and $\theta_{-i}^{(2)}$, such that agent i can cause all four outcomes to happen by changing its own bid alone (although we are considering the semi-shattering notion so the order in which they happen does not matter).

Let the weighted sum of the bids of the agents other than i for the first (second) profile under outcome A be a_1 (a_2), under outcome B be b_1 (b_2), and so on. Also, let the weights on agent i 's bid under outcomes A through D in the the GSP outcome function (Equation 3.1) be α_A through α_D . (Note that the premise of our theorem implies that $\alpha_A > \alpha_C > \alpha_D > \alpha_B$.)

Let us assume (without loss of generality) that $b_1 - a_1 < b_2 - a_2$ and that A will happen against $\theta_{-i}^{(1)}$ and B will happen against $\theta_{-i}^{(2)}$ (if the inequality does not hold, we can reverse the labels on the θ_{-i} 's). In order to cause A to happen against the first opponent profile and B against the second, the following inequalities must hold (we assume that ties are broken consistently so that an agent cannot use them to semi-shatter).

$$A \text{ happens against } 1 \begin{cases} \alpha_A \theta_i + a_1 > \alpha_B \theta_i + b_1 \\ \alpha_A \theta_i + a_1 > \alpha_C \theta_i + c_1 \\ \alpha_A \theta_i + a_1 > \alpha_D \theta_i + d_1 \end{cases}$$

$$B \text{ happens against } 2 \begin{cases} \alpha_B \theta_i + b_2 > \alpha_A \theta_i + a_2 \\ \alpha_B \theta_i + b_2 > \alpha_C \theta_i + c_2 \\ \alpha_B \theta_i + b_2 > \alpha_D \theta_i + d_2 \end{cases}$$

than higher ranking ones. We use 10 to the negative power since it is easy to conceptualize.

By simplifying the above equations we derive the following set of constraints.

$$\frac{c_1 - a_1}{\alpha_A - \alpha_C} < \theta_i < \frac{b_2 - d_2}{\alpha_D - \alpha_B}$$

$$\frac{d_1 - a_1}{\alpha_A - \alpha_D} < \theta_i < \frac{b_2 - c_2}{\alpha_C - \alpha_B}$$

In order to semi-shatter C and D with C happening against the first set of bids by the other agents and D against the second we have the following inequalities generated in the same fashion.

$$\frac{b_2 - d_2}{\alpha_D - \alpha_B} < \theta_i < \frac{c_1 - a_1}{\alpha_A - \alpha_C}$$

In order to semi-shatter over C and D in the opposite direction (with D first and C second), the constraints would change to the following.

$$\frac{b_2 - c_2}{\alpha_C - \alpha_B} < \theta_i < \frac{d_1 - a_1}{\alpha_A - \alpha_D}$$

Our assumption that agent i could semi-shatter both sets of outcomes when the other agents have more than a single type leads to a contradiction since the two sets of inequalities cannot be simultaneously satisfied. \square

This result, in conjunction with Theorem 2 from Chapter 2, implies that under some preference distributions the efficiency bound for the GSP is arbitrarily inefficient, and, since it is an upper bound, the inefficiency exists under any strategy profile.

Corollary 7. *For any setting there exists a distribution over agent preferences such that the upper bound on expected efficiency (Equation 2.2) for the GSP mechanism's outcome function is arbitrarily less than fully efficient.*

3.4 The Premium GSP mechanism

To address GSP's inexpressiveness without making the mechanism much more cumbersome, we introduce a new mechanism that only slightly increases the expressiveness. Later we

show empirically that this slight increase is extremely important in that it removes most of the efficiency loss entailed by GSP’s inexpressiveness in many realistic settings.

The new mechanism separates the positions into two classes: premium and standard, and each agent can submit a separate bid for each class. We call this the *premium generalized second price (PGSP)* mechanism. The premium class might contain, for example, only the top position—as in our experiments.

The premium position(s) are assigned as if a traditional GSP were run on the premium bids (the top premium position goes to the agent with the highest premium bid, etc.). The standard positions are then assigned among the remaining agents according the traditional GSP mechanism run on their standard bids.

3.5 Computing the efficiency bound

The results in Section 3.3 prove that there exist distributions over agent preferences for which the GSP is arbitrarily inefficient. However, in order to measure the inefficiency in practice we must be able to compute the value of the efficiency bound for a given distribution over agent preferences. In this section, we describe two general techniques for doing that. They take as input a distribution over agent preferences with a finite number of types (this distribution could be learned from data or approximated by a domain expert) and provide the value of the upper bound on the mechanism’s most efficient equilibrium. Although we present our techniques in the context of ad auctions, they can easily be generalized for use in other domains.

3.5.1 Integer programming formulation

First, we will describe an integer programming formulation for computing the bound. The program includes a binary decision variable, z_o^t , for each outcome, o , and each joint type of the agents, t . A value of 1 for z_o^t denotes that outcome o will be chosen by the mechanism when the agents have the joint type t , a value of 0 indicates that the outcome will not be chosen under t . The program also includes continuous variables representing the agents’ expressions (bids in the context of sponsored search) under each of their types, $\theta_i^{t_i}$. (We limit

these expressions to be between 0 and 1, without loss of generality.)⁴ The following objective function is used to maximize the expected efficiency of the mechanism. To accomplish this, we sum the welfare over all types and outcomes, weighted by their probabilities.

$$(3.2) \quad \max_{z_o^t, \theta_i^{t_i}} \sum_{t \in T^n} P(t) \sum_{o \in O} z_o^t W(t, o)$$

The first set of constraints enforces that exactly one outcome is chosen for each joint type. There are $|T^n|$ such constraints.

$$(3.3) \quad \text{s.t. } (\forall t \in T^n) \quad \sum_{o \in O} z_o^t = 1$$

The next set of constraints ensures that for each z_o^t variable that is set to 1, the agents' expressions under type t do indeed cause the outcome function to choose outcome o . This set includes one constraint for each joint type and each pair of distinct outcomes. Thus there are $|T^n| \times (|O|^2 - |O|)$ such inequality constraints.⁵ These constraints depend on the outcome function of the mechanism we are studying. For GSP's outcome function, the constraints are as follows (we use M to denote a sufficiently large number such that the sum of all the agents' expressions cannot exceed it):

$$(3.4) \quad (\forall t, \forall o, \forall o' \neq o) \quad \sum_i (\theta_i^{t_i} 10^{-R(i,o)}) > \sum_i (\theta_i^{t_i} 10^{-R(i,o')}) - (1 - z_o^t)M$$

Finally, we have constraints on the decision variables:

$$(3.5) \quad (\forall t, \forall o) z_o^t \in \{0, 1\}, (\forall i, \forall t_i) 0 \leq \theta_i^{t_i} \leq 1$$

An ad auction with k positions and n agents with two types each has $\frac{n!}{(n-k)!}$ distinct outcomes and 2^n joint types. The integer program has $|O| \times |T^n|$ binary decision variables, making it prohibitively large for general purpose integer program solvers, such as CPLEX, for mechanisms with more than 3 agents. However, these solvers do not explicitly take advantage of certain aspects of the problem structure, for example the fact that only one outcome can be chosen for each joint type.

⁴In practice, the expression space would have to be discrete as well (e.g., discretized to accommodate a currency), however we assume that such a discretization would always be possible at a fine enough level so as not to affect our simulations. This makes the search problem easier as well, since it allows us to use linear programming to assess the feasibility of an outcome assignment.

⁵In practice, we ensure that these inequality constraints are strict by adding a small ϵ term to one side.

3.5.2 Tree search for computing the bound

To address this problem, we developed a general tree search technique based on A* for computing the bound. We have applied the technique to GSP and PGSP on instances with up to five agents to find provable inefficiency. (In this chapter, we only report results with up to four agents in order to provide a larger number of experiments.)⁶

Each level of the search tree corresponds to a different joint type. Each branch corresponds to the assignment of an outcome to the joint type. The tree has maximum depth $|T^n|$ and branching factor $|O|$. Figure 3.1 illustrates the search tree.




Figure 3.1: Part of the search tree for a distribution with 3 types, $[t_1^n, t_2^n, t_3^n]$, and 4 outcomes $[A, B, C, D]$. Circles represent internal nodes and squares represent leaf nodes. The dashed nodes are not expanded, but they would be considered by the algorithm. The expanded path corresponds to the assignment of $[A, C, C]$ to types t_1^n , t_2^n , and t_3^n , respectively.

⁶While five-agent instances may seem particularly small for some keyword auctions, the agents in our simulations can also be thought of as representing segments or blocks of agents that all behave the same way. Under this assumption, all of our analysis regarding overall effects on efficiency would still hold true.

At any node j a partial assignment of outcomes to joint types can be constructed by traversing the edges from j to the root. We will denote the set of all joint types in the partial assignment at node j as T_j^n . For each type $t_j^n \in T_j^n$ we will denote the outcome it is assigned under the partial assignment at node j as o_{t_j} . In addition, for each joint type t^n we will denote any one of the outcomes that maximize social welfare as o_t^* (i.e., $o_t^* = \arg \max_o W(t^n, o)$).

As usual, our search orders the nodes in its open queue according to an admissible (i.e., optimistic) heuristic. We developed a custom upper-bounding heuristic for the search algorithm, which enables early pruning of branches, and thus dramatically reduces total search time, while preserving optimality of the search algorithm. The heuristic approximates the expected efficiency of the best assignment originating from a particular node under the assumption that any unassigned types will be assigned optimally.⁷ The priority of a node j , $\tilde{f}(j)$, is given by the expected welfare of its current partial assignment plus the expected welfare of the optimal assignment for any unassigned types:

$$(3.6) \quad \tilde{f}(j) = \sum_{t_j \in T_j^n} P(t_j)W(t_j, o_{t_j}) + \sum_{t \notin T_j^n} P(t)W(t, o_t^*)$$

Interesting aspects of our upper bounding heuristic include that 1) it can be applied for any mechanism regardless of its expressiveness (and it is, in a sense, the only nontrivial such heuristic), and 2) much of the computation can be pre-calculated and cached before the search.

The $\tilde{f}(j)$ approximation is guaranteed to be greater than or equal to the true optimal value of any feasible assignment that descends from node j . It may overestimate this value if the optimal assignment is not achievable due to inexpressiveness, but it has the benefit of serving as a valid upper bound on the expected efficiency achievable by the mechanism. By using the A* node selection strategy, our search ensures that any node that it visits has a lower (or equal) \tilde{f} value than any previously visited node. Thus, the \tilde{f} value of the current node is a continually tightening upper bound on the mechanism's expected efficiency, and it can be provided at any time during the search. In our experiments we were occasionally forced to terminate the search early in order to evaluate a greater number of preference

⁷We need only calculate o^* once at the beginning of the search. It can be reused later by removing outcomes that are assigned.

distributions. In these cases we reported the \tilde{f} value of the last feasible node that was visited as our upper bound.

Whenever a node is popped off the front of the open queue, its feasibility is checked. In both types of ad auction mechanisms we study, this check involves solving a linear feasibility problem (LFP). The LFP has a set of constraints similar to those described in Equation 3.4, however the assignment of outcomes to types is fixed and there are no binary decision variables. If the node is not feasible, its children are not placed on the open queue. Specifically, at any node j we verify that there exist expressions for the agents conditioned on their types, $\theta_i^{t_i}$, which satisfy the following constraints.

$$(3.7) \quad (\forall t_j \in T_j^n, \forall o' \neq o \in O) \quad \sum_i \left(\theta_i^{t_i} 10^{-R(i, o_{t_j})} \right) > \sum_i \left(\theta_i^{t_i} 10^{-R(i, o')} \right)$$

3.6 Experiments with GSP and PGSP

In this section, we discuss the results of experiments using our search technique to compute the upper bound for the GSP mechanism and the slightly more expressive PGSP mechanism.

In order to gain additional insight, we also discuss the performance of the two mechanisms when agents use the straightforward strategy of always bidding their valuation for the top position (in the PGSP they bid their valuation for the top premium position and the top non-premium position as their two bids). We call the resulting efficiency *GSP heuristic* and *PGSP heuristic*, respectively. Such a heuristic, or a variation of it that would not affect the rankings (e.g., with bids shaded by a constant amount), is likely to be used in practice and provides a useful baseline to compare with the value of the cooperative equilibrium.

Our experiments consist of collections of runs, each involving randomly generated instances with different parameter settings. The parameters are chosen to investigate circumstances under which the inexpressiveness of the GSP mechanism is costly (i.e., when the upper bound is low) and when it is not. Each instance in one of our experiments represents a single auction for a single keyword with three or four agents.

3.6.1 Experimental setup

Based on recent work examining different advertising attitudes on the Internet, in our experiments each agent is either a *brand advertiser* (with probability p_B) or a *value advertiser* (with probability $1 - p_B$) [15]. Brand advertisers always prefer higher positions over lower ones. A value advertiser generally does not prefer the highest positions because middle positions tend to have higher *conversion rates* (e.g., the user’s probability of buying something conditional on having clicked is higher). Others have also begun to explore the implications of value advertisers [29], and some work has even used the exact distributions we developed [141]. There has also been some recent work on click models to support this experimental setup [56, 117]. Figure 3.2 illustrates prototypical brand and value preferences over different positions based on their rank.




Figure 3.2: Example of prototypical valuations for brand and value advertisers. The brand advertiser shown has $\mu = 1$ and the value advertiser has $\mu = 0.5$ (as defined in Table 3.1). Valuations are shown in expectation, not per click (these values will be negative when the amortized cost per click of running the site is high or the expected value of a conversion is low, which we vary in our experiments). Rank 0% means the bottom position and Rank 100% means the top position.

Random instance generation

When we generate instances for our experiments we assume an agent’s valuation for being assigned a particular position is the expected value of having its ad displayed in that position. We now describe how we generate preferences for brand and value advertisers. Let “clk” denote the event that the ad is clicked, and “cnv” denote that the click results in a conversion (e.g., a sale or user registration). Let C_i denote the amortized cost per click of running agent i ’s web site, and $V_i(\text{cnv})$ be the expected value of a conversion to agent i . Then, the expected value to agent i of having an ad in position ranked R is given by the following.

$$(3.8) \quad E[V_i(R)] = P(\text{clk}|R, i) [P(\text{cnv}|\text{clk}, R, i)V_i(\text{cnv}) - C_i]$$

In order to keep the experiments simple, and to focus on the impact of expressiveness, we assume that agents in the same instance are relatively similar. For one, we assume that the marginal cost of a click $C_i = C = \$1$ for all agents.⁸ Unless otherwise specified, we assume that $V_i(\text{cnv}) = V(\text{cnv}) = \50 for all agents. We assume that $P(\text{cnv}|\text{clk}, i) = P(\text{cnv}|\text{clk}) = 10\%$ for all agents (note that this probability is not what differentiates the two type of advertisers, but rather the probability that a conversion comes from a given rank). We also assume that click-through rates conditional on the rank of an ad’s position are the same for all agents. The specific rates are given in Table 3.1, along with the default values for all parameters. These click-through rates are from an Atlas Institute Digital Marketing publication [38]. They were also used by Even-Dar *et al.* in their experiments [60].

Rather than generating arbitrary values of $P(\text{cnv}|\text{clk}, R, i)$, we assume that the probability of a conversion coming from a particular rank, $P(R|\text{cnv}, i)$, is normally distributed. The mean, μ , of this distribution is randomly chosen from $[0, 1]$ for each agent, once for the case where she is a brand advertiser and once for the case where she is a value advertiser. (We also normalize the value of R to be between 0 and 1, so that, for example, the third position out of four has rank 0.25.) Values of μ closer to 1 indicate that the agent’s conversions tend to come from higher ranked ads, those closer to 0 indicate that conversions tend to come from lower ranked ads. The values of μ for the brand and value advertisers are given in Table 3.1, unless otherwise specified.

⁸This may seem like a large value for some settings, however, since we consider only the fraction of optimal efficiency achieved by each mechanism, this cost is only important in relation to the value of a conversion, which we vary widely in our experiments.

Parameter	Default value
$P(\text{clk} R = 1)$	10%
$P(\text{clk} R = 2)$	7.74%
$P(\text{clk} R = 3)$	6.66%
$P(\text{clk} R = 4)$	5.74%
$P(\text{cnv} \text{clk})$	10%
p_B	50%
$C(\text{clk})$	\$1
Brand μ	$\sim \text{Uniform} [.8, 1]$
Brand σ	25% of μ
Value μ	$\sim \text{Uniform} [.4, .6]$
Value σ	25% of μ
$V_i(\text{cnv})$	\$35 to \$150

Table 3.1: Default settings for each parameter in our instance generation model.

We transform $P(R|\text{cnv}, i)$ into $P(\text{cnv}|\text{clk}, R, i)$ using Bayes' rule (and the observation that the cnv event implies the clk event):

$$(3.9) \quad P(\text{cnv}|\text{clk}, R, i) \propto P(R|\text{cnv}, i)P(\text{cnv}|\text{clk}, i)$$

Each data point in each figure below is the average over 50 instances. The confidence intervals represent standard error. (They are often so tight that they are barely visible.)

3.6.2 Experiment 1: Varying agents' profit margin

In our first set of results we vary the expected value of a conversion, $V(\text{cnv})$, between \$35 and \$150 (i.e., 35 to 150 times the cost per click of running the site). The results are shown in Figure 3.3 and Figure 3.4. The values are reported in terms of the percentage of the optimal efficiency achievable.

These results demonstrate that when conversions generate relatively low profits, the efficiency loss due to inexpressiveness in the GSP mechanism, as measured by the upper




Figure 3.3: The value of the upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on four-agent instances. Results are averaged over 50 runs with different expected values for a conversion.




Figure 3.4: The value of the upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on three-agent instances. Results are averaged over 50 runs with different expected values for a conversion.

bound, is more than 30%. As the profit margin of the agents increases, this loss decreases to around 10%.

Additionally, the results show that the efficiency bound for the slightly more expressive PGSP mechanism is nearly 100% in all cases. This suggests that the added expressiveness in the PGSP is well suited to capture all the different types of preferences we generated.

We also see that the efficiency of the heuristic bidding strategy follows a similar qualitative pattern to the upper bound, which lends additional support to our findings. Since this heuristic strategy represents a natural strategy that is likely to be taken by many agents in this domain, our results with this strategy suggest that 1) the bound is meaningful in describing the efficiency of the mechanism, and 2) the conclusions apply more broadly than for fully rational, game-theoretic agents.

The instances with three and four agents exhibit relatively similar values for the efficiency bound at each value of $V(\text{cnv})$ when all other parameters are held at their default values. (The slightly higher values of the bound for the four agent instances with larger conversion values can be partially explained by the fact that around 25% of these instances were terminated early due to our 20 minute timeout, however these timeouts were distributed evenly throughout the parameter space).

3.6.3 Experiment 2: Varying agent diversity

The second experiment examines how the loss due to inexpressiveness depends on how similar value advertisers are to brand advertisers. Specifically, we vary the position that generates the most value for value advertisers. (Brand advertisers still always prefer the highest position the most.) In each run the mean of $P(R|\text{cnv}, i)$ for each value advertiser is drawn uniformly from an interval of size 0.2 (i.e., $\mu \sim \text{Uniform}[a, a + 0.2]$). The results are shown in Figure 3.5 and Figure 3.6. The x-axis indicates the mid-point of the interval used in each run, which is also the expected value of μ for each value advertiser.

These results demonstrate that the need for additional expressiveness is greatest when the value advertisers prefer middle ranking positions, as is typically the case in practice. For example, when those agents prefer the middle rank, the GSP can achieve at most 85% efficiency (with the heuristic bidding strategy achieving less than 75%) on average, whereas




Figure 3.5: The value of our upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on four-agent instances. Large values of $E[\mu]$ correspond to runs in which higher ranking positions are more valuable for the value advertisers and vice versa.

the PGSP can achieve over 95% for the bound (with the heuristic bidding strategy achieving about 85%). The expressiveness is less crucial when the value advertisers are more akin to the brand advertisers (i.e., large $E[\mu]$) or when they drastically differ (i.e., small $E[\mu]$).

Again, the efficiency of the heuristic bidding strategy follows a similar qualitative pattern to the upper bound. Also, our results on instances with fewer agents show that the cost of inexpressiveness tends to be more severe when the GSP mechanism is run with four agents than when it is run with three, suggesting that these issues may be magnified as the number of agents increase.

3.7 Conclusions and future research

In this chapter we operationalized our theoretical framework from Chapter 2 by developing a methodology for comparing mechanisms with different degrees and forms of expressiveness and applied it to sponsored search.




Figure 3.6: The value of our upper bound on expected efficiency and the efficiency of the heuristic bidding strategy for the GSP and PGSP mechanisms on three-agent instances. Large values of $E[\mu]$ correspond to runs in which higher ranking positions are more valuable for the value advertisers and vice versa.

We began by proving that for some preference distributions the most commonly used sponsored search mechanism, GSP, is arbitrarily inefficient. In order to measure the inefficiency in practice we developed a general tree search technique for computing an upper bound on a mechanism’s expected efficiency. We concluded with a series of experiments comparing the GSP to our slightly more expressive mechanism, PGSP, which solicits an extra bid for premium ad positions. We generated a range of realistic preference distributions, based on published industry knowledge, and applied our search technique to compare the efficiency bounds for the two mechanisms. We also examined the performance of the mechanisms when agents use a straightforward heuristic bidding strategy.

Our results suggest that the GSP’s efficiency loss due to inexpressiveness can be dramatic. It is greatest in the practical case where some agents (“brand advertisers”) prefer top positions while others (“value advertisers”) prefer middle positions. The loss is also worst when agents have small profit margins. Despite the fact that our PGSP mechanism is only slightly more expressive (and thus not much more cumbersome), it removes almost all of the

efficiency loss in all of the settings we study.

One future research opportunity involves using real data for ads placed in different positions. (Such data sets have recently been made publicly available.) However, due to the difficulty in obtaining data about preferences and conversions, it will likely be necessary to adapt our methodology to incorporate other meaningful ways of measuring the inefficiency. For example, rather than relying on real preference data to entirely replace the simulated distributions, one will likely need to develop a “hybrid” distribution that is still partially simulated, but is more directly informed by real-world data than those we described. One method for doing this would involve positing a parametric distribution for advertiser valuations (e.g., similar to the models we proposed) and clustering advertisers based on bids into the two types of advertisers. The parameters for for each type of advertiser could then be inferred from bids or conversion data and the methodology we described could be applied using the resulting models.

It would also be interesting to consider how other types of expressiveness could benefit the GSP. For example, expressions that allow advertisers to bid higher for certain types of users that are likely to convert (e.g., “premium” users). This could result in greater efficiency than even what is possible using the fully expressive mechanism in our experiments since that mechanism, as we described it, does not allow such expressions.

Another future direction is to adapt recent methods for computing equilibria in sponsored search mechanisms by modeling them as action graph games [91, 141] to compute equilibria for our PGSP mechanism. One can then compare the equilibria under the PGSP and GSP mechanisms in terms of revenue and efficiency to see if they match the results of our cooperative bound and heuristic. This analysis can be performed under a variety of preference distributions to determine the types of preferences for which the PGSP is more efficient and profitable in equilibrium.

The methodology we have developed can also be adapted and extended to other application domains, such as combinatorial auctions and voting mechanisms. For combinatorial auctions, we have already adapted our theoretical framework to channel-based mechanisms, which provide an abstraction of almost all commonly studied auction mechanisms. One can operationalize our theory further in that domain by developing search algorithms that automatically design channel-based mechanisms subject to limits on the number of channels

given to the agents.

Chapter 4

Expressiveness in Privacy Mechanisms

4.1 Introduction

The past few years have seen an explosion in the range of websites allowing individuals to exchange personal information and content that they have created. These sites include location-sharing services, which are the focus of this chapter, social-networking services, and photo- and video-sharing services. While there is clearly a demand for users to share this information with each other, there is also substantial demand for greater control over the conditions under which this information is shared. This has led to expanded privacy and security controls on some services, such as Facebook, but designers of others appear reluctant to make this change. One reason for this reluctance may be that more complex privacy controls typically lead to more complex and hard-to-use interfaces. What is missing is a methodology for determining the relative importance of different expression types for a given user population.

In this chapter, we begin by applying our theoretical framework for studying expressiveness to the domain of privacy. We define a class of mechanisms that we call *privacy mechanisms*, or mechanisms that allow individuals to control the circumstances under which certain pieces of private information are shared. In this domain, our adapted notions of expressiveness can be used to characterize the level of control an individual has over how his or her private information is released. Using our theoretical framework, we prove that more expressiveness can be used to design more efficient privacy mechanisms – or mechanisms that allow individuals to share more of the information they want to share, without violating their privacy preferences.

Next, using our theoretical framework as a foundation, we proceed to describe how the benefits of expressiveness for privacy mechanisms can be quantified in practice for location-sharing privacy mechanisms. Around one hundred different location-sharing applications exist today [143]. These applications allow users to share their location (frequently, their exact location on a map) and other types of information, but have extremely limited privacy mechanisms. Typically, they only allow users to specify a *white list*, or a list of individuals with whom they would be willing to share their locations at any time [143]. Despite the number of these types of applications available, there does not seem to be any service that has seen widespread usage. One possible explanation for this slow adoption has been established by a number of recent papers, which demonstrate that individuals are concerned about

privacy in this domain [40, 53, 54, 80, 90, 122, 144]. However, our work is the first, to our knowledge, to study location-privacy preferences at a detailed enough level to address the question of whether or not more expressive privacy mechanisms may help alleviate these concerns.

We present the results from a user study where we tracked the locations of 27 subjects over three weeks in order to collect their stated location-privacy preferences in detail. Each day, for each of the locations a subject visited, we asked whether or not he or she would have been willing to share that location with each of four different groups: close friends and family, Facebook friends, the university community, and advertisers.¹ Throughout the study, we collected more than 7,500 hours of location information and corresponding privacy preferences. In contrast to some earlier research that identified the requester’s identity [53] and user’s activity [52] as primarily defining privacy preferences for location sharing, we find that there are a number of other critical dimensions in these preferences, including time of day, day of week, and exact location.

We characterize the complexity of our subjects’ preferences by measuring the accuracy of different privacy mechanisms with different levels and types of expressiveness. We consider privacy mechanisms that allow a user to share his or her location based on the group of the requester, the time of day of the request, whether or not the request is made on a weekend, and his or her location at the time of the request. Using the detailed preferences we collected during the location tracking phase, we identify each subject’s most accurate collection of rules,² or *policy*, under each type of privacy mechanism. To test the effectiveness of the different mechanisms, we measure the accuracy with which each is able to capture our subjects’ preferences,³ while varying assumptions about the relative cost of revealing a private location, and about our subjects’ tolerance for user burden. Our accuracy metric is equivalent to the expected efficiency of a privacy mechanism where agents have policy-based utility functions, which we will define in Section 4.2.

As one might expect, we find that more complex expression types, such as those that

¹In this study, we do not account for different usage levels of Facebook (e.g., by considering the number of friends of the users). However, we believe this is an interesting issue to consider in future work.

²A rule is defined naturally for each type of privacy mechanism, e.g., a span of time, or rectangle enclosing multiple locations.

³The notion of accuracy we use in this chapter is also equivalent to the expected efficiency of the privacy mechanism under certain reasonable assumptions about the users’ utility functions.

allow users to specify both location- and time-based rules, are more accurate at capturing the preferences of our subjects under a wide variety of assumptions. More surprising is the magnitude of accuracy improvement — in some cases more complex expression types can result in almost three times the average accuracy of white lists. White lists appear to be particularly ineffective at capturing our subjects’ preferences. Even relatively simple extensions, such as those that allow rules based only on time of day, can yield a 33% increase in average accuracy, assuming that our subjects are privacy sensitive. This finding is also consistent with results from our pre-study survey, where subjects reported being significantly more comfortable with the prospect of sharing their location using time- and location-based rules, compared to white lists.

In addition to accuracy, we measure the amount of time each day that our subjects would have shared their location under each of the different privacy mechanisms. Interestingly, we find that more accurate privacy mechanisms also lead to more sharing. This result, which at first may seem counter intuitive, actually makes sense: when users have complex privacy preferences and are given limited settings, they generally tend to err on the safe side, which causes them to share less.⁴ This may explain why some social networking sites, such as Facebook, have begun to move toward more expressive privacy mechanisms — if users end up sharing more, the services are more valuable. The lack of sharing we observe with simple privacy mechanisms may also help explain the slow adoption of today’s location sharing applications.

While our results suggest that more expressive privacy mechanisms are necessary to capture the true location-privacy preferences of the user population represented by our subjects, these mechanisms do not come without a cost. More complex expression types generally imply additional user burden, especially if they require users to specify significantly more rules than their simple counterparts. To address this, we examine a number of different privacy mechanisms, which range from being fairly simple to more complex, under varied assumptions regarding the amount of effort our subjects would be willing to exert while creating their policies. For the purposes of this chapter, we use the number of rules a policy contains as a proxy for the user burden involved in specifying it. Our findings suggest that, while limiting policies to a small number of rules dampens the accuracy benefits of expressive

⁴Another way to think about this is that privacy-sensitive users first attempt to find rules that minimize mistaken sharing, and among those possible rules choose the ones that maximize the amount of time shared.

privacy mechanisms, they generally remain substantially more accurate than white lists.

The user study presented in this chapter also demonstrates a general methodology for characterizing the tradeoffs between more expressive privacy mechanisms and accuracy (or efficiency) in a number of privacy and security domains. At a high level, the methodology involves i) collecting highly detailed preferences from a particular user population, ii) identifying policies for each subject under a variety of different privacy or security mechanisms, and iii) comparing the accuracy of the resulting policies under a variety of assumptions about the sensitivity of the information and tolerance for user burden.

The rest of this chapter proceeds as follows. In the next section, we present a discussion of the theoretical background behind our study. In Section 4.3, we provide the details of the methods used in conducting our user study and analyzing the data. In Section 4.4, we present a detailed analysis of our data. Finally, we present some conclusions and possibilities for future work in Section 4.5.

4.2 Theoretical background

One key difference between the formal model of expressiveness in this chapter, and that of our other work is a move to a single agent setting. In this chapter, we assume that the behaviors of agents other than the one making an expression are stochastic, rather than strategic (e.g., requests for one’s private information are assumed to come from some probability distribution, rather than the behavior of other rational agents). Despite this difference, we will show that our theoretical framework for studying expressiveness can be naturally applied to this domain.

4.2.1 A general privacy mechanism model

The formal setting we study in this chapter is that of a single request for a piece of private information, such as an individual’s geographical location. We assume that a request can be described by a vector of m attributes, $\vec{a} = \{a_1, a_2, \dots, a_m\}$, such as the individual behind the request, or the time the request was placed. In general, each of these attributes can be discrete valued or real valued (however, in practice we discretize real-valued attributes, such

as time). We assume that the attribute vector, \vec{a} , of a request is stochastically drawn from the set of all possible requests, \vec{A} , according to a joint probability distribution, which we denote as $P(\vec{a})$.

In our model, an agent interacting with the mechanism has a type, t , which is unknown to the mechanism. The agent's type is drawn according to some probability distribution, $P(t)$, from the set of all possible types, T , and represents the agent's attitude toward releasing any piece of private information under any circumstance (the set of all types can be finite or infinite). For example, an agent may have a type that is highly secretive about releasing its location during certain times of day, or its type may be more concerned about releasing certain locations.

The agent interacts with the mechanism by making an expression about its privacy preferences, which we denote as θ , from the space of all possible expressions, Θ . Based on the privacy preferences that the agent expresses and the attributes of a request, the mechanism computes the value of a binary outcome function, $f(\Theta, \vec{A}) \rightarrow \{0, 1\}$. The outcome function determines whether the request is granted (i.e., when $f(\theta, \vec{a}) = 1$) or denied (i.e., when $f(\theta, \vec{a}) = 0$).⁵ In our model, the piece of private information under consideration to be shared (e.g., a user's location) is considered to be a fixed value that is outside the scope of the mechanism (i.e., it is not given as an argument to the mechanism). However, we assume that the mechanism has access to that information to use aspects of it when determining an outcome (e.g., when a user expresses a location-based rule, we assume the mechanism can lookup the user's last known location for any incoming request).

We assume that the agent has a utility function, u , which depends on the agent's type, the attributes of a request, and the outcome chosen by the mechanism. The utility function maps these inputs to a real-valued utility indicating how happy or unhappy the agent is with the outcome chosen by the mechanism, $u(T, \vec{A}, \{0, 1\}) \rightarrow \mathbb{R}$. We will also define an agent's strategy, $h(T) \rightarrow \Theta$, as a mapping from each possible type to an expression. A strategy dictates how the agent will interact with the mechanism depending on its type. Typically we assume that the agent will choose a strategy, h^* , that maximizes its expected utility.⁶

⁵In this chapter, we assume that the outcome function is binary: it either grants or denies a request. However, it is possible to generalize our notion of binary outcomes to include cases where a request can be granted to differing degrees, such as releasing an individual's city, rather than exact location.

⁶Note that when a user has a highly negative utility associated with mistakenly revealing a piece of

$$h^*(t) = \arg \max_{\theta} \int_{\vec{a}} P(\vec{a}) u(t, \vec{a}, f(\theta, \vec{a}))$$

Using this model, we can describe the expected efficiency of a particular privacy mechanism with the following equation (where expectation is taken over the possible types of the agent and the different possible request attributes, when attributes and types are considered to be discrete the integrals in the following equation would be summations instead), which is similar to Equation 2.1:

$$(4.1) \quad E[\mathcal{E}(f)] = \int_t P(t) \int_{\vec{a}} P(\vec{a}) u(t, \vec{a}, f(h^*(t), \vec{a}))$$

4.2.2 Policy-based utility functions

In our empirical analysis we focus on one simple class of utility functions, which we call *policy-based utility functions*. An agent always has some underlying privacy preference function, $\pi(T, \vec{A}) \rightarrow \{0, 1\}$, which indicates the outcome that the agent prefers for any possible request. With a policy-based utility function we assume that the agent suffers a cost c whenever the mechanism inappropriately grants a request, the agent suffers a cost of c' whenever the mechanism denies a request that should have been granted, and the agent receives reward r whenever the mechanism correctly releases information. Typically we assume that the cost for mistakenly revealing a piece of private information is much greater than the reward for correctly sharing it, (i.e., $c \gg r$). Table 4.1 illustrates this class of utility functions under each of the four possible scenarios: i) the mechanism correctly grants, ii) correctly denies, iii) inappropriately grants or iv) inappropriately denies.

4.2.3 Expressiveness and efficiency in privacy mechanisms

We will now demonstrate that a privacy mechanism's expected efficiency is closely related to its expressiveness level. Our first result shows that when designing a privacy mechanism, any

information, this maximization becomes a maximization over the expressions that minimize the likelihood of that occurrence.

	Mechanism deny ($f(\theta, \vec{a}) = 0$)	Mechanism allow ($f(\theta, \vec{a}) = 1$)
Agent deny ($\pi(t, \vec{a}) = 0$)	$u(t, \vec{a}, f(\theta, \vec{a})) = 0$	$u(t, \vec{a}, f(\theta, \vec{a})) = -c$
Agent allow ($\pi(t, \vec{a}) = 1$)	$u(t, \vec{a}, f(\theta, \vec{a})) = -c'$	$u(t, \vec{a}, f(\theta, \vec{a})) = r$

Table 4.1: An illustration of the policy-based utility function class under each of the four possible scenarios: i) the mechanism correctly grants, ii) correctly denies, iii) inappropriately grants or iv) inappropriately denies.

increase in allowed expressiveness can be used to achieve strictly higher expected efficiency.⁷

Theorem 8. *For any utility function, distribution over agent types, and distribution over request attributes, the expected efficiency (given in equation 4.1) for the best privacy mechanism limiting an agent to impact dimension d increases strictly monotonically as d goes from 1 to d^* , where d^* is the minimum impact dimension needed to reach full efficiency.*

Proof. The set of mechanisms with impact dimension d is a super-set of the mechanisms with impact dimension $d' < d$. Thus the fact that the efficiency for the best mechanism increases weakly monotonically is trivially true. The challenge is proving the strictness of the monotonicity.

Consider increasing d from $d^{(1)} < d^*$ to $d^{(2)} > d^{(1)}$. Let $G^{(1)}$ be the best set of impact vectors that an agent could distinguish between when restricted to $d^{(1)}$ vectors (i.e., the set of impact vectors that would maximize the mechanism's expected efficiency). We know that there are at least $d^* - d^{(1)} \geq 1$ impact vectors needed to reach full efficiency that cannot be expressed, and thus at least that many impact vectors that are absent from $G^{(1)}$. When we increase our expressiveness limit from $d^{(1)}$ to $d^{(2)}$, we can add one of those missing vectors to $G^{(1)}$ to get $G^{(2)}$. Since $G^{(2)}$ allows an agent to distinguish among all the same vectors as $G^{(1)}$ and an additional vector that corresponds to a more efficient set of outcomes, the new mechanism with impact dimension $d^{(2)}$ has a strictly higher expected efficiency. \square

In addition, we see that even a small increase in allowed expressiveness can be used to achieve an arbitrarily large increase in a mechanism's expected efficiency.

⁷The results in this section have been adapted to this domain from the results in Chapter 2. The primary departure from that work is the move to a stochastic setting, rather than a strategic setting.

Theorem 9. *There exists a utility function, a distribution over types, and a distribution over request attributes such that the best privacy mechanism limited to impact dimension d is arbitrarily less efficient than that of the best privacy mechanism limited to impact dimension $d + 1 < d^*$, where d^* is the minimum impact dimension needed for full efficiency.*

Proof. Since an agent’s utility function can depend arbitrarily on its type and the attributes of a request, we can construct a scenario in which the agent requires impact dimension at least $d + 1$ or it will experience an arbitrarily high cost. First we must ensure that the agent has at least $d + 1$ types with non-zero probability. Next we choose a set of impact vectors, $G^{(1)}$, of size $d + 1$. For each of the distinct impact vectors in $G^{(1)}$ we can ensure that it gives the agent arbitrarily more utility than all other impact vectors for at least one of the agent’s types. By the pigeon hole principle, the agent will be unable to express at least one of the impact vectors in $G^{(1)}$ in any mechanism with impact dimension d . Thus increasing a limit on impact dimension from d to $d + 1$ will lead to an arbitrary increase in efficiency. \square

These results taken together suggest that privacy mechanisms can be made significantly more efficient by designing them with greater levels of expressiveness. Throughout the rest of this chapter, we will describe an extensive user study that we performed to test these findings in practice.

4.3 Methods for our user study

We will now discuss the methods used to conduct and analyze our location sharing user study. We provide an overview of our study, details of the software we used to conduct it, descriptions of the privacy mechanisms we consider, and a description of the methods we use to analyze them. Our study also serves as a methodology for quantifying the benefits of different types of privacy mechanisms in a wide variety of domains. At a high level, the methodology proceeds as follows: First, we collect highly detailed privacy preferences from our subjects. Next, we identify different privacy mechanisms with varying levels and forms of expressiveness. We then identify a policy for each subject under each privacy mechanism, while taking into consideration various levels of user burden. Finally, we compare the accuracy of the policies for different privacy mechanisms under a variety of assumptions about the sensitivity of the information and tolerance for user burden.

4.3.1 Study overview

The data for our study was collected over the course of three weeks in early November 2009. We supplied 27 participants with Nokia N95 cell phones⁸ for the entire study. Each subject was required to transfer his or her SIM card to the phone we provided and use it as a primary phone at all times. This requirement ensured that subjects kept the phones on their person, and charged, as much as possible. Each of the phones was equipped with our location-tracking program, which recorded the phone's location at all times using a combination of GPS and Wi-Fi-based positioning.

Each day, subjects were required to visit our web site where the locations recorded by their phones were filtered into distinct location observations. For each location a subject visited, we asked whether or not he or she would have been comfortable sharing the location at that time with different groups of individuals and advertisers. These groups consisted of close friends and family, Facebook friends, people within the university community, and advertisers. While no location sharing to others actually occurred, we solicited the names of people from the different groups (other than advertisers) so that the questions the subjects answered were more meaningful. We later displayed these names in each audit question presented to the subject, as shown in Figure 4.2. For the Facebook group, we automatically scraped the names of all of our subjects' friends and presented them with a random selection in each audit.

We also administered surveys before and after the study to screen for participants, measure the level of concern about privacy that people had about sharing their location information, and collect relevant demographics. The full text of these surveys can be found in an appendix at the end of this chapter. The screening process ensured subjects had, or were willing to purchase, a cellular data plan with a compatible provider.

Subjects were paid a total of \$50-\$60, corresponding to \$30 for their successful participation in the study, and \$20-\$30 to reimburse them for the data plan that was required by the location-tracking software.

⁸These phones were generously provided by Nokia.

4.3.2 Software

The primary materials we used in our experiment included location-tracking software written for the Nokia N95 phone and a web application that allowed subjects to audit their location information each day.

Location-tracking software

Our location-tracking software is written in C++ for Nokia's Symbian operating system. It runs continuously in the background, and starts automatically when the phone is turned on. During normal operation, the software is completely transparent – it does not require any input or interaction. When designing our software, we faced two primary challenges: i) managing its energy consumption to ensure acceptable battery life during normal usage, and ii) determining the phone's location when indoors or out of view of a GPS signal. To address these challenges, our software is broken down into two modules: a *positioning module* that tracks the phone's location using a combination of GPS and Wi-Fi-based positioning, and a *management module* that turns the positioning module on and off to save energy.

Positioning module. To estimate the position of the phone, our positioning module makes use of the Nokia N95's built in GPS, and Wi-Fi units. When activated, the positioning module registers itself to receive updates from the GPS unit at a regular interval (15 seconds). When the GPS unit is able to determine the phone's position, the positioning module records its latitude and longitude readings. Whenever the positioning module is active it also records the MAC addresses and signal strengths of all nearby Wi-Fi access points at a regular interval (3 minutes). We are able to use this information to determine the physical address of the phone with a service called Skyhook Wireless.⁹ While the positioning module is active, it sends all location information to our server using the phone's cellular data connection in real time.

Management module. Our initial tests revealed that leaving the GPS unit on continuously resulted in an unacceptable battery life of 5-7 hours on average. The management

⁹Details about the Skyhook API are available at <http://skyhookwireless.com/>.

module uses the N95's built in accelerometer to address the issue of energy consumption. It constantly monitors this low energy sensor, and only activates the positioning module when the accelerometer reports substantial motion. In practice we found that this improved the phone's battery life to 10-15 hours on average.¹⁰

Web application

Each day, subjects were required to visit our web site to audit the locations they visited that day. The locations were first filtered, then presented to the subjects to audit.

Location filtering. When a subject logs into our web application, it iterates through each of the GPS and Wi-Fi readings that have been recorded since the last time the user audited his or her locations. Each of these readings is either aggregated into a location observation, if the user stood still, or a path observation, if the user moved.¹¹ A new location observation is created when a subject has moved more than 250 meters from his or her last known location and remained stationary again for at least 15 minutes.

Audit administration. After a subject's locations have been filtered, our web application takes the subject through a series of pages that trace his or her new locations in chronological order. Each page displays a location on a map, inside a 250-meter ring, indicating the subject's estimated location during a particular time period. The times when the subject arrived and departed from the location are indicated next to the map. Each page also includes a link that allows subjects to report that an observation was completely inaccurate (inaccurate observations accounted for about 2% of the time, and are removed during our analysis). A screen shot of the user interface for this part of the web application is shown in Figure 4.1.

Underneath the map, our web application presents four questions, each corresponding to a different group of individuals. Figure 4.2 shows an example screen shot of a question for

¹⁰For more details about this process, see the description of a similar technique used by Wang *et al.* for managing energy consumption while tracking users with mobile devices [153].

¹¹Path observations between locations were also depicted on some pages. However, we do not address those observations here since they accounted for less than 1% of the observed time.

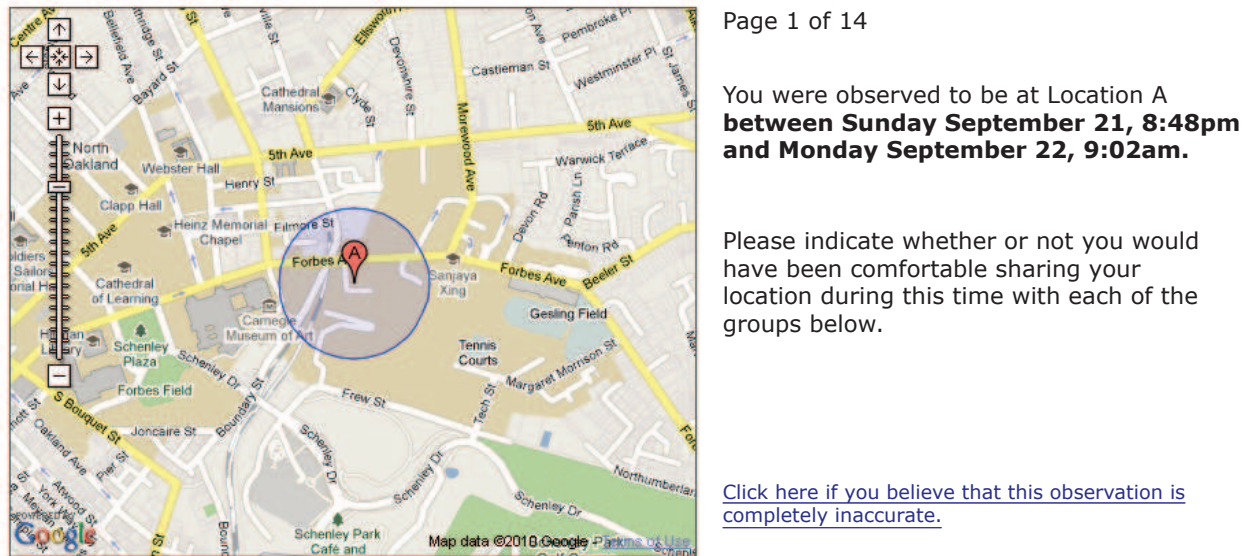
the friends and family group. Each question asks whether or not the subject would have been comfortable sharing his or her location with the individuals in one of the groups. The groups we asked about in our study were: i) close friends and family, ii) Facebook friends, iii) anyone associated with our university, and iv) advertisers. Subjects are given the option of indicating that they would have shared their location during the entire time span indicated on the page, none of the time span, or part of the time span (when part of the time is chosen, a drop down menu appears allowing the subjects to specify which part of the time they would have allowed, as shown in Figure 4.2). Questions about the friends and family and Facebook groups include a fourth option, allowing subjects to indicate that they would have been comfortable sharing their location with some of the individuals in the group, but not all of them.¹²

4.3.3 Privacy mechanisms we compare

In our analysis (Section 4.4.3), we focus on evaluating the accuracy of the following different privacy mechanisms, which range from being fairly simple to more complex. We will illustrate the differences between them by considering a hypothetical user named “Alice,” who wishes to share her location only with her friends when she is at home, on the weekends, between the hours of 9am and 5pm. In the absence of a rule that explicitly shares one’s location, we assume that the default behavior of a sharing service would be to deny.

- **White list.** White lists are the least expressive privacy mechanism we consider. They only allow users to indicate whether or not they would be comfortable sharing their location with each group at all times and locations. The accuracy of white lists can be viewed as a measure of the importance of a requester’s identity in capturing users’ privacy preferences. White lists are user friendly, since they only require a single rule indicating who can view one’s location.

¹²The partial group option was chosen about 20% of the time for Facebook friends. However, 89% of the time this option was chosen by a subject, the subject also reported that he or she would have been comfortable sharing with either friends and family, or the university community. These subjects were most likely considering one or both of these two groups as subgroups of Facebook friends. This hypothesis is further supported by the fact that 82% of the subjects reported in the post-study survey that they did not feel there were any relevant groups missing from our list. For these reasons, we treat this response as denying the entire group in our subsequent analysis.




Would you have been comfortable sharing your location **between Sunday September 21, 8:48pm and Monday September 22, 9:02am** with:

Figure 4.1: A screen shot of our web application displaying an example location on a map between 8:48pm and 9:02am.

Using a white list, our hypothetical user, Alice, would need to indicate *who* (individually or by group) is allowed to see her location. Similarly, she may also create a rule that everyone is allowed to see her at all times with a list of exceptions (i.e., a black list). Alice’s policy under this mechanism would not match her preferences, since friends on her white list would be able to see her anytime and anywhere.

- **Location (Loc).** The Loc mechanism allow users to indicate specific locations that they would be comfortable sharing with each group. This mechanism is more expressive than a white list, since it can be used like a white list by sharing all locations with a group. The accuracy of Loc can be seen as a measure of the importance of location in capturing users’ privacy preferences. A single location rule is defined by a latitude-longitude (lat-lon) rectangle and a set of people or groups who can view the user’s location within the rectangle.¹³

¹³It is also reasonable to assume that a single Loc rule involves only a single location, rather than a lat-lon



Your Close Friends and Family?
(e.g., Jim, Mary, Pam, etc.)

Yes, during this entire time

No, not during any of this time

Yes, during part of this time...

Yes, for some of these people

I would have been comfortable sharing my location from:

9/21 8: 48 pm

to:

9/22 9: 02 am

[Add an additional time span.](#)

Figure 4.2: A screen shot of an audit question asking whether or not a subject would have been comfortable sharing the location displayed on the map with the friends and family group. An audit question, like the one shown here, appeared below the map for each of the groups, at each location a subject visited. Drop down menus are only displayed because “Yes, during part of this time...” is selected.

Alice would need to create a rule allowing her friends to view her location when she is at home, by indicating it with a rectangle on a map, but this policy would not match her preferences precisely, since her friends could see whether or not she was home at night or on a weekday.

- **Time.** The Time mechanism allows users to indicate time intervals (discretized into half-hour blocks) during which they would be comfortable sharing their locations with

rectangle. However, we chose to consider an entire rectangle as a single rule since the service that our study is modeled on, Locaccino, allows that.

each group (this mechanism does not consider the day of the week). Similar to Loc, Time is more expressive than a white list, since white listing for an individual or group can be simulated by granting them access at all times. The accuracy of Time can be seen as a measure of the importance of the time of day in capturing users' privacy preferences. For some distributions over possible requests, the Time mechanism is more expressive than the Loc mechanism, but for other distributions the opposite is true. In other words, neither the Loc mechanism nor the Time mechanism is more expressive for all possible request distributions. A single time rule is defined by a start time, an end time, and a set of people or groups who can view the user's location between the two times.

Under the Time mechanism, Alice would need to create a rule sharing her location with her friends between 9am and 5pm, regardless of where she was and the day of week. Alternatively, she could err on the safe side and choose to share a smaller time window during which she feels she is more likely to be home. In either case, Alice's policy would not match her preferences, since her friends could potentially see her location when she was somewhere other than at home.

- **Time with weekends (Time+).** The Time+ mechanism is the same as Time, but it allows users to indicate time intervals that apply only to weekdays, only to weekends, or to both. Thus, it is also more expressive than Time. The improvement in accuracy of Time+ over Time can be viewed as the importance of weekends in capturing our subjects' privacy preferences. A single rule under Time+ is defined by a start time, an end time, a flag indicating whether it applies to weekdays, weekends, or both, and a set of people or groups who can view the user's location, between the two times, on the specified type of day.

Under the Time+ mechanism, Alice would need to create a rule sharing her location with her friends, between 9am and 5pm on weekends only, regardless of where she was. As with Time, Alice's policy would not match her preferences, since her friends could see her location when she was somewhere other than at home, but with Time+ this could not happen on a weekday.

- **Location and time (Loc/Time).** The Loc/Time mechanism combines the Loc and Time expression types described above and is, thus, more expressive than those two

mechanisms. Loc/Time allows users to indicate time intervals during which they would be comfortable sharing specific locations with each group. The accuracy improvement of Loc/Time over Loc and Time individually can be viewed as the importance of offering both types of expressions together. A single Loc/Time rule is defined by a start time, an end time, a lat-lon rectangle, and a set of people or groups who can view the user’s location when he or she is within the rectangle between the two times.

Under the Loc/Time mechanism, Alice would need to create a rule allowing her friends to see her when she is at home, from 9am to 5pm, regardless of the day of week. In this case, Alice’s policy would not match her preferences, since her friends could potentially see her at home on a weekday.

- **Location and time with weekends (Loc/Time+).** Loc/Time+ is the same as Loc/Time, but it allows users to indicate time intervals that apply only to weekdays, only to weekends, or to both. This is the most expressive privacy mechanism we consider.

Under Loc/Time+, Alice would be able to express her true privacy preferences with a single rule: allow her friends to see her when she is at home, from 9am to 5pm, on weekends only.

4.3.4 Measuring accuracy with variable cost

In order to measure the accuracy of different privacy mechanisms, we first identify a collection of rules, or a *policy*, for each subject, under each of the different mechanisms described in Section 4.3.3. For a subject, i , a privacy policy, p , and group, g , we define the accuracy of the policy for i and g using two functions, `correct_hrs` and `incorrect_hrs`. The functions take as input i , p , and g , and return the number of hours correctly shared and incorrectly shared, respectively, by subject i , with group g , under p . These statistics are easily computed from our data for any possible policy, since we can simulate what the policy would have done at each of the locations a subject visited, and compare that to their stated preferences for that location. We normalize the accuracy to be a fraction of the time shared by each subject’s *optimal policy*, or the policy that perfectly matches the subject’s preferences (i.e., shares whenever the subject indicated he or she would do so, and does not share at any other times

or locations).

In our analysis, we will consider the accuracy of different privacy mechanisms while varying assumptions about our subjects' tolerance for mistakes. For this, we define a penalty term, or cost, c , associated with mistakenly revealing a piece of private information. In our analysis, we vary c from 1 to 100 and investigate the impact it has on accuracy and sharing under the different privacy mechanisms. Varying c amounts to varying the ratio between the reward for revealing a location when a subject indicated that he or she would have shared it and the penalty for revealing it when he or she indicated not being comfortable with having it shared. At the lowest level (when $c = 1$) these two occurrences are equally rewarded and penalized, respectively. When $c = 100$, mistakenly revealing a location is considered to be one-hundred times as bad as correctly revealing it. This level of cost is essentially equivalent to the assumption that our subjects would be very cautious, and never make policies that mistakenly revealed their locations. Varying this cost helps to account for differences between subjects and across potential applications.¹⁴ Accuracy for a policy, group, and subject is given by the following equation, where p^* is the subject's optimal policy.

$$(4.2) \quad \frac{\text{correct_hrs}(i, p, G) - c \times \text{incorrect_hrs}(i, p, G)}{\text{correct_hrs}(i, p^*, G)}$$

The accuracy of the best policy for any subject, group, and privacy mechanism, will always be between zero and one. It can never be below zero, because an empty policy achieves zero accuracy, and it can never be above one, since we normalize the accuracy for each subject using the accuracy of the best possible policy for that subject.¹⁵ Note that the average accuracy value we report is equivalent to the expected efficiency of each mechanism, as defined in Section 4.2, assuming that subjects have policy-based utility functions and are equally likely to receive requests at all times. The utility functions would provide a reward of

¹⁴We assume that there is no penalty for mistakenly withholding a location, since our post-study survey results suggest that subjects had relatively little dis-utility at this prospect. However, this can easily be added as an additional cost to the accuracy calculation in Equation 4.2.

¹⁵When a subject indicated that he or she would never have shared their location with a particular group, thereby making the accuracy equation undefined, we report the accuracy for that subject and group as one, since we assume that the default behavior of the system is to deny access, which is consistent with the subject's preferences.

$r = 1$ unit per hour for a location is correctly shared (i.e., given to a group during a time that was marked as allowed). We assume that the subjects would receive 0 utility whenever their locations are blocked (i.e., $c' = 0$), rather than penalizing them for any missed opportunities, and subjects pay a cost c whenever their locations are inappropriately shared (i.e., shared with a group during a time that was marked as not allowed). Our results consider several different utility functions by varying the value of c .

4.3.5 Identifying privacy policies with user-burden considerations

In Section 4.4.3, we consider how accurate the different privacy mechanisms are under the most accurate policy for each subject with no rule limit. Then, we consider the effect of limiting the number of rules to account for user-burden tolerance. In both cases, the accuracy values that we report can be taken as upper bounds on the accuracy we would expect in practice, since subjects may not always create the most accurate possible policy.

With no rule limit, a subject's most accurate policy for a given group and privacy mechanism can be easily computed by identifying all possible *atomic rules* for the group and mechanism (e.g., rules that apply only to a single location, or a single half-hour block). We then greedily add an atomic rule whenever it would result in positive accuracy for the subject (i.e., when it is correct more than $1/c$ of the time). This is guaranteed to identify the most accurate policy, since the search decomposes in the following straightforward way: each group, time, location and location/time pair can be allowed or disallowed independently (when rules regarding weekends and weekdays are considered, we treat times on the two types of days independently). For example, the effect on overall accuracy of adding a rule sharing a particular location does not depend on which other locations the policy ends up sharing.

Like many other combinatorial problems (e.g., knapsack, job-shop scheduling, graph coloring), the problem of identifying the most accurate policy for a given subject and privacy mechanism becomes substantially harder with a limited resource, such as rules. For example, with a limit on the number of rules the greedy solution is no longer guaranteed to identify the most accurate policy. To address this problem, we developed a tree-search technique, based on the well-known A* search algorithm, for computing a subject's most accurate policy with no more than k rules.

Each level of the search tree corresponds to one of the rules in the policy, and each branch represents a particular rule that can be included. For example, one branch could correspond to the rule “University community and Friends can see me at any location, between 8:00am and 7:00pm, on weekdays.” Thus, at any node, j , with depth d , a policy with d rules can be constructed by traversing the edges from j to the root. Figure 4.3 illustrates part of a search tree using the Loc/Time+ mechanism.

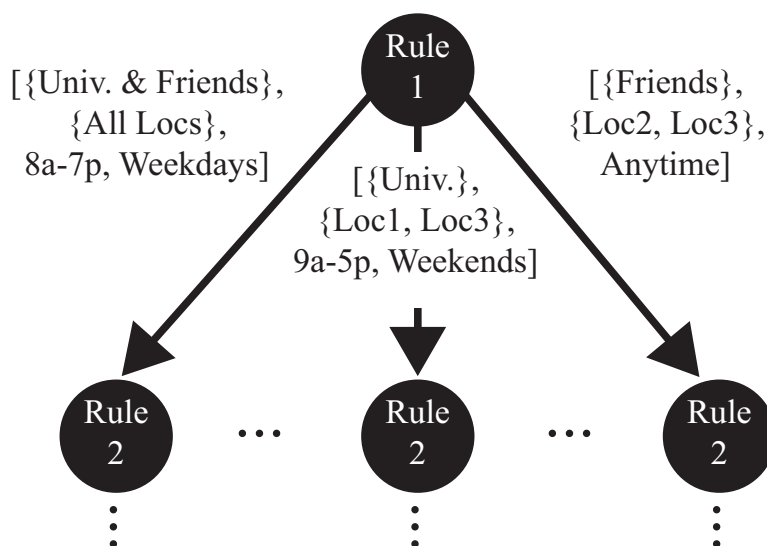


Figure 4.3: Part of a search tree for identifying a subject’s most accurate privacy policy using the Loc/Time+ mechanism.

Our search begins at the root node, and constructs one child node for each of the possible rules a user could add, given the type of expressions available. The nodes are added to a priority queue, called the *open queue*. Nodes are then popped off the open queue one at a time until a leaf node (i.e., node with depth k) is reached. Whenever a node, j , is removed from the open queue, a child of j is added to the queue for each of the remaining *feasible* rules. A rule is considered feasible for inclusion in children of j if it does not *overlap* with any rule that is already in the policy represented by j . Two rules overlap if they refer to the same place, time, or place and time, for Loc, Time (Time+), and Loc/Time (Loc/Time+), respectively.

As usual, our search orders the nodes in its open queue according to an admissible (i.e.,

optimistic) heuristic. The heuristic approximates the accuracy of any policy with k rules originating from a particular node as the total accuracy of the rules included so far, plus the accuracy of a greedy solution over the remaining feasible rules with no rule limit. In our case, this technique of using a greedy solution with a relaxed constraint as a heuristic is guaranteed to produce a solution with greater than or equal to the best total accuracy of any set of k rules descending from node j . However, it may overestimate this value if the greedy solution uses more than k rules. By using the A* node selection strategy, our search ensures that any node it visits has a lower (or equal) accuracy than any previously visited node, thus making the first depth- k solution reached provably the most accurate one possible.

If we were to consider every possible atomic rule at each level of this search tree it would be intractable for the more expressive privacy mechanisms. There are 48 different 30-minute spans in a day, each span can apply to weekdays, weekends, or both, and subjects visited about 10 locations on average. If we assume that any possible combination of locations can be grouped together (this is an overestimate because some combinations will be infeasible) the tree would have $2^{10} \times 48 \times 3 = 147,456$ nodes at each level, and more than 10^{20} nodes in total with four rules. To address this, we also losslessly compress the search space by preprocessing each subject's ground truth policy according to the following technique. For Loc rules, individual locations are grouped together into complex locations if they are audited the same way at all times (i.e., sharing them always results in positive accuracy for the same groups) and it would be possible to draw a rectangle around them without including any of the subject's other locations. For Time (and Time+) rules, individual half-hour spans are grouped together if they are audited the same way every day (and type of day for Time+). For Loc/Time (Loc/Time+) rules, locations are grouped together if they are always audited the same way based on time of day and it would be possible to draw a rectangle around them without including any other locations. With these preprocessing steps in place, we can identify policies for each subject, and privacy mechanism, typically in a matter of seconds.

4.4 Empirical findings from our user study

Before we present our analysis on measuring the effects of different privacy mechanisms, we will describe our survey findings, the general mobility patterns we observed, and some high-

level statistics that demonstrate the complexity of our subjects' location-privacy preferences.

For all statistical tests of significance, we use two-sample independent t-tests with unequal variances, unless otherwise noted. Throughout this section, we report p values of less than 0.05 as significant and less than 0.1 as marginally significant. Due to the large number of quantities we compare in our analysis, in most cases we also present 95% confidence intervals on estimates assuming the underlying data is normally distributed.¹⁶

4.4.1 Survey results

Our 27 subjects were all students or staff at our university. The sample was composed of 73% males with an average age of about 22 years old. Undergraduates made up 58% of our sample, graduate students made up 35%, and two people (7%) were staff members.

In our pre-study survey, we asked participants about how comfortable they would be if close friends and immediate family, Facebook friends, members of the university community, or advertisers could view their locations at anytime, at times they had specified, or at locations they had specified. Based on ratings on a 7-point Likert scale (ranging from “not comfortable at all” to “fully comfortable”), we found that, in general, participants were more comfortable with their close friends and family locating them than their Facebook friends, people within their university community, or advertisers.

Within each group, we found that respondents had relatively equal levels of comfort for time-based or location-based rules (the differences were not statistically significant). However, it is interesting to note that location had a substantially higher average score than time for the advertiser group, since we later find that this is the only group for which the difference between the accuracies of Loc and Time mechanisms is marginally significant. The average scores for this question are shown in Table 4.2.

We also found that subjects reported that they would be significantly more comfortable, on average, for the Facebook friends, university community, and advertiser groups, using location- and time-based rules than with white lists. For example, for the advertisers group, our subjects indicated that they would not be comfortable if their locations were shared all

¹⁶We present these confidence intervals in lieu of an explicit ANOVA test, which is primarily used to accommodate experiments with different groups of people and no clear ordering.

Group	Anytime	Location	Time
Friends and family	5.00	6.08	6.36
Facebook friends	3.64	4.88	5.40
University community	3.28	4.56	5.00
Advertisers	2.60	4.32	3.60

Table 4.2: The average report on our pre-study survey of how comfortable subjects would have been on a 7-point Likert scale from “not comfortable at all” to “fully comfortable” if their location could be checked by each of the groups “Anytime,” “At locations you have specified,” or “At times you have specified.”

the time ($M=2.6$); but at times ($M=3.60$) or locations ($M=4.32$) they had specified, their comfort levels would significantly increase.

After completing our study, we asked our participants how bad they thought it would have been, on a 7-point Likert scale from “not bad at all” to “very, very bad,” if the system had shared their information at times when they did not want it to be shared, or if the system had withheld their location when they wanted it to be shared. Table 4.3 shows the average report for each type of mistake and each group.

Group	Mistakenly withheld	Mistakenly revealed
Friends and family	3.00	3.26
Facebook friends	2.30	3.70
University community	2.07	4.26
Advertisers	1.67	4.74

Table 4.3: The average report of how bad subjects thought it would have been, on a 7-point Likert scale from “not bad at all” to “very, very bad,” if their location were mistakenly withheld from or revealed to each of the groups.

Our subjects reported significant levels of dis-utility at the prospect of their locations being mistakenly shared with the university community, Facebook friends, and advertisers groups, with the worst being advertisers, where 33% of the participants chose 7 on the scale and 50% choose 5 or more. In contrast, our subjects reported relatively little dis-

utility at the prospect of their locations being mistakenly withheld. We also see an inverse relationship between the average report within groups, such that groups where mistakenly revealing is worse tend to have lower reports for mistakenly withholding. This lends support to the hypothesis that our subjects would tend to share less when given less expressive privacy mechanisms, since they report being far more concerned with inadvertent disclosure of their location than with it being withheld, on average. It also supports the assumption we make later that the cost associated with accidental disclosure is much larger than the cost associated with accidental withholding.

We asked our subjects how often they would have answered the questions differently if we had actually been sharing their locations. The majority of subjects (about 70%) responded that they would have rarely or never answered differently. Another 15% said they would have answered differently some of the time, and the rest said most or all of the time.

4.4.2 Mobility patterns and preference statistics

On average, our subjects were observed for just over 60% of the time during our experiment, and our observations were distributed relatively evenly throughout the day. We found that, on average, subjects would have been comfortable sharing their locations about 93% of the time with friends and family, 60% of the time with Facebook friends, 57% of the time with university community, and 36% of the time with advertisers.

Figure 4.4 shows how our subjects' preferences varied with time of day and day of week. It shows the average percentage of time subjects were willing to share during each half-hour interval separately for weekdays and weekends.

Preferences for the friends and family group are largely unaffected by time of day or day of week. However, the results show substantial variation in preferences based on time of day and day of week, for the other three groups. For these groups, we see almost twice as much sharing during the day on weekdays as at night and on weekends. On weekends we also see slightly greater preferences for sharing during the evening.

About half of our subjects visited 9 or fewer distinct locations throughout the study, and 89% visited 14 or fewer (the max was 27, the min was 3). A subject was considered to have visited a distinct location only if it was visited for at least 15 minutes, and was at least 250




Figure 4.4: The average percentage of time shared with each group during each thirty-minute interval throughout the day on weekdays (top) and weekends (bottom).

meters from all other locations that the subject visited.

We found that, on average, subjects spent significantly more time at one location than any other (most likely their homes). We also found that the time spent at a location appeared to drop off significantly for the second, third, fourth and fifth most visited locations. Table 4.4 shows the average percentage of time a subject spent at his or her five most visited locations, and the average percentage of time that he or she would have shared that location with each

of the groups. On average, our subjects were more willing to share their second most visited location than their first. For university community and advertisers they were willing to share it almost twice as often. This suggests that this was most likely a more public location, such as somewhere on or near the university campus [142].

Location rank (time spent)	Time spent	Time shared w/ group			
		FF	FB	UC	AD
1st	66%	93%	58%	48%	29%
2nd	20%	94%	65%	77%	55%
3rd	6%	90%	61%	62%	41%
4th	3%	99%	55%	61%	35%
5th	1%	97%	48%	52%	35%

Table 4.4: The average percentage of time a subject spent at his or her five most visited locations, and the average percentage of time he or she would have shared that location with friends and family (FF), Facebook friends (FB), university community (UC), and advertisers (AD).

These results suggest mobility patterns similar to those observed by Gonzalez *et al.*, who found that human trajectories tend to be very patterned, with people visiting a small number of highly frequented places [63]. These results also help explain our later finding that the Loc mechanism only requires a few rules to realize most of its benefits.

4.4.3 Measuring the effects of different privacy mechanisms

We will now present analysis quantifying the relative effects of different privacy mechanisms, in terms of accuracy and amount of time shared. We consider the results statistically, and under a wide range of assumptions, including varying levels of user burden.

The relative accuracy scores of the different privacy mechanisms provide quantitative measures of their importance for capturing our subjects' preferences. Consequently, the differences between the accuracy of different privacy mechanisms measures the importance to our subjects of the preference dimensions on which they differ. Under some circumstances, we find substantial accuracy benefits from more expressive privacy mechanisms, lending

support to the conclusion that the location sharing preferences revealed by our study are fairly rich. We also find that the more accurate policies typically result in subjects sharing more, not less, of their information, since users tend to err on the safe side when the cost of mistakenly revealing their information is relatively high.¹⁷

Results regarding policy accuracy

Our first set of results, presented in Figure 4.5, investigates the accuracy of each of the different privacy mechanisms, for each of the groups we asked about. For these results, we hold the cost of mistakenly revealing a location to be fixed at $c = 20$, which is equivalent to assuming that subjects view mistakenly revealing their location as twenty times worse than correctly sharing. We highlight our results for this value of c based on the post-study survey results presented in Table 4.3, which showed that subjects were significantly concerned with mistakenly revealing their location to each of the groups other than their close friends and family. Our next set of results will consider varying this cost to account for differences between subjects and groups.

Our first observation is that, with $c = 20$, none of the privacy mechanisms we consider are able to achieve 100% accuracy for any of the groups. Even the accuracy of the most accurate mechanism and group, Loc/Time+ for friends and family, is significantly less than 100% (as evidenced by the fact that its 95% confidence interval ends below that point). This demonstrates that a non-trivial subset of our subjects had preferences that alternated between sharing and hiding the same location, at the same time, on different days of the week (most likely due to other contextual factors).

With $c = 20$, the average accuracy of the different privacy mechanisms has a wide range across groups, from about 28% (white lists for advertisers) to 88% (Loc/Time+ for friends and family). There is also a moderately large range in accuracy, across groups, for the same simple mechanisms (e.g., white lists range from 28% to 68%). However, the range across

¹⁷When we report the average percentage of time shared here, we include locations that would have been shared by the mechanism even if the subjects indicated they would not have wanted to share them. When we consider only locations that the subjects wanted to share, we see the same general pattern in all results and for most c values the difference between the sharing statistics calculated these two different ways is negligible.




Figure 4.5: The average accuracy (bars indicate 95% confidence intervals) for each group, under each of the different privacy mechanisms. For these results, we hold constant the cost for inappropriately revealing a location at $c = 20$.

groups is substantially smaller for more expressive mechanisms (e.g., *Loc/Time+* ranges from 68% to 88%). This suggests that expressive privacy mechanisms mitigate the importance of a requester’s identity in capturing our subjects’ preferences.

The range of average accuracies within groups is smaller, but still substantial. For example, within the advertisers group, accuracies range from 68%, for *Loc/Time+*, to 28%, for white lists. For the Facebook friends and university community groups, we also observe a more than two times increase in accuracy of *Loc/Time+* over white lists. The fact that such ranges in accuracy exist within groups further demonstrates that our subjects had diverse privacy preferences that could not all be captured simply by the requester’s identity.

For advertisers, the expressive mechanisms (i.e., *Loc/Time* and *Loc/Time+*) are significantly more accurate than white lists, the *Time* mechanism, and the *Time+* mechanism. *Loc* alone is also significantly better than white lists, and marginally significantly better than *Time*. The relative importance of location-based rules for this group is consistent with our pre-study survey findings presented in Table 4.2.

In other groups, we see statistical ties between *Loc*, *Time+*, and *Time*, although *Loc*

tends to be the best of the three on average. We also see that the mechanisms allowing users to distinguish between weekdays and weekends can offer substantial benefits over their simpler counterparts (e.g., for university community Time+ is about 15% more accurate than Time), but these differences are typically not statistically significant.

For university community and Facebook friends, we find that Loc/Time+ is significantly more accurate than all of the mechanisms other than Loc/Time. For university community, we find that Loc/Time is significantly more accurate than white lists, Time, and Time+, and marginally significantly more accurate than Loc. For Facebook friends the finding is nearly the same, but Time+ is statistically tied with Loc/Time. This demonstrates the importance of weekends in capturing our subjects' preferences about sharing their location with Facebook friends.

All of these results taken together suggest that, with $c = 20$, our subjects could expect significant accuracy improvements from more expressive privacy mechanisms, and further confirms the hypothesis that the privacy preferences revealed by our study are complex.

Our next set of results, shown in Figure 4.6, investigates the impact of varying the cost associated with mistakenly revealing a location, for the Facebook friends group. We present these results for Facebook friends only because we believe that this group is of general interest, and results for other groups were qualitatively similar.

These results demonstrate that the accuracy benefits of more expressive privacy mechanisms are greatest when information is more sensitive. For example, when $c = 1$, we find that there are no statistically significant differences between any of the mechanisms. In this case, the difference between the most expressive mechanism, Loc/Time+, and the simplest, white lists, is only marginally significant. However, the accuracies of less expressive mechanisms drop steeply as the cost of inappropriately revealing one's location increases. For example, the accuracy of white lists drops from 61% at $c = 1$, to almost half of that, or 34%, at $c = 25$, and drops to 28% by the time we reach $c = 100$. Similar patterns are seen with all of the less expressive mechanisms, such as Time, Time+, and Loc. This drop is due to the fact that, as this cost goes up, the policies we identify are more restrictive (e.g., by concealing more often). Thus, they provide lower accuracy because they have missed more opportunities to share.

Each of the mechanisms also reaches a plateau at different values of c . The plateau occurs




Figure 4.6: The average accuracy for the Facebook friends group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.

when the subjects have been forced to hide as much as they can, and only reveal times or locations that are never private. The accuracies of more expressive mechanisms, such as Loc/Time and Loc/Time+, deteriorate far less, far slower, and with plateaus beginning at far lower costs than simple types (e.g., the plateau for Loc/Time+ begins at $c = 10$, whereas white lists continue to lose accuracy throughout the entire range). This demonstrates how more expressive privacy mechanisms can add substantial value for privacy-sensitive users.

Results regarding amount of time shared

We now consider how the policies we identified for different privacy mechanisms effect the amount of time our subjects would have shared with each of the groups. Figure 4.7 shows the average percentage of time that each subject would have shared, under each of the different mechanisms, with a fixed cost of $c = 20$ for mistakenly revealing a location.

Here we see results similar to those in Figure 4.5, such that more accurate policies also tend lead to more sharing with each group. For example, for the Facebook friends, university community, and advertiser groups, we see about twice as much sharing with Loc/Time+ versus white lists, and in each case this difference is statistically significant (the difference




Figure 4.7: The average percentage of time shared (bars indicate 95% confidence intervals) with each group under each of the different privacy mechanisms. For these results, we hold constant the cost for inappropriately revealing a location at $c = 20$.

between Loc/Time and white lists in each case is also marginally significant). It is also interesting to note that Loc and Time+, which are relatively simple, still result in substantial increases in sharing over white lists for the advertiser group (19% and 17% vs. 10%, respectively); however, neither of these differences is statistically significant.

That sharing increases with more accurate privacy mechanisms is explained by the fact that, when $c = 20$, mistakenly revealing one’s location is substantially worse than mistakenly withholding it. This, in turn, leads to policies that tend to err on the safe side and share less.

Our next set of results, presented in Figure 4.8, considers the effect of varying the cost of mistakenly revealing a location on the amount of time shared under each privacy mechanism. Again, we limit our presentation to the Facebook friends group, since results for other groups were qualitatively the same.

The findings here are similar to those presented for accuracy in Figure 4.6, with a few notable differences. We see a general trend from more to less sharing as c increases, with plateaus beginning at around $c = 10$, however the plateaus are far more dramatic and jagged than with accuracy. This is because we only observe effects on sharing when individual rules




Figure 4.8: The average percentage of time shared with the Facebook friends group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.

are made more restrictive, rather than the smooth descent in accuracy that leads to the restriction.

As with accuracy, the decline in sharing with more expressive privacy mechanisms, such as *Loc/Time+* and *Loc/Time*, is less steep, and slower than that of the less expressive ones. A higher value for c represents the assumption that users are more concerned about privacy. Thus, this demonstrates how it can actually be in a service's best interest to offer more expressive privacy mechanisms, in order to increase contributions from privacy-sensitive users.

One final take away from this analysis is the magnitude of the increase in sharing with highly privacy-sensitive users, under the most expressive privacy mechanism, *Loc/Time+*, versus white lists. For $c = 100$, which corresponds to the assumption that users will make policies that never give out private information, we see a more than three and a half times increase in the average percentage of time shared with the Facebook friends group.

All of these results taken together suggest, somewhat counter-intuitively, that offering richer privacy settings may, in fact, make good business sense, since it will result in privacy-sensitive users sharing more information.

Results under user-burden considerations

In practice, we do not expect users to necessarily specify the most accurate policy matching their preferences, especially under the more expressive privacy mechanisms, such as Loc/Time+, where user interfaces can be cumbersome. To test the effects of such user-burden considerations on our conclusions, we analyze the effect of limiting the number of rules in policies for each of the privacy mechanisms.

Our first set of results under user-burden considerations is presented in the four panels of Figure 4.9, one for each group. It shows the accuracy of each mechanism, while varying a limit on the number of rules from one to five or more. This set of results is modeled after a scenario where sharing one’s location with all four groups is possible within a single application, and users specify rules that apply to combinations of these groups. We operationalize this by identifying the most accurate policy with a *global rule limit*, rather than a limit that applies to each group individually. For each of the different privacy mechanisms, we identify policies that equally weight accuracy among the groups. In other words, results shown in the four panels for a global rule limit of two amounts to finding the best policy with only two rules when it comes to sharing with all of the four groups.

Unsurprisingly, we find that tighter rule limits generally dampen the accuracy benefits of more expressive privacy mechanisms. Yet, we see that Loc/Time+ and Loc/Time have substantial benefits, in terms of average global accuracy, with as few as one or two rules. For example, if we consider the global average accuracy across all groups, with only a single rule we already see a marginally significant benefit from Loc/Time+ (51%) over white lists (35%). With two rules, the difference between the accuracy of Loc/Time+ (54%) and white lists is significant, and the difference between the accuracy of Loc/Time (50%) and white lists is marginally significant. This demonstrates how more expressive privacy mechanisms can be better than less expressive ones at capturing the preferences of our subjects, while requiring only a small number of rules.

When we examine the effects of a global rule limit on the accuracies within individual groups, rather than the global average accuracy, with two rules we find a significant accuracy improvement for the university community group from Loc/Time+ (52%) over white lists (31%), and a marginally significant difference between those two mechanisms for advertisers (45% vs. 28%). With three rules, the difference in accuracy between Loc (49%) and white




Figure 4.9: The average accuracy (vertical axis) achieved by each of the different privacy mechanisms, for each of the different groups, varying a global limit on the number of rules (horizontal axis, from one to five or more) in a policy. We hold constant the cost for inappropriately revealing a location at $c = 20$, and identify policies with the highest possible total accuracy across all groups, while weighting each group equally.

lists is significant, and the difference between Loc and Time (33%) is marginally significant. Interestingly, with three rules, the Loc/Time and Loc/Time+ mechanisms actually perform worse for advertisers than the less expressive Loc mechanism. This is because under the more expressive mechanisms, the three rules are primarily being used to achieve greater accuracy in other groups, whereas the accuracy of Loc tends to plateau with two rules. This plateau can be explained, in part, by the general mobility patterns presented in Table 4.4, which show that subjects tended to spend about 80% of their time at two distinct locations.

Our final set of results, presented in Figure 4.10, is modeled after a service where users can share locations with a single group only, such as all of one’s Facebook friends. Here we limit the rules that apply to a group individually, rather than imposing a global limit. We present the results for the Facebook friends group only, but results for other groups were similar.




Figure 4.10: The average accuracy (bars indicate 95% confidence intervals) achieved by each of the different privacy mechanisms for the Facebook friends group, while varying a limit on the number of rules in a policy that apply to Facebook friends only. We hold constant the cost for inappropriately revealing a location at $c = 20$

By comparing the results in Figure 4.10 to those in the top right panel of Figure 4.9, we find that with an individual rule limit the accuracy benefits of more expressive privacy mechanisms are realized with fewer rules. For example, we find that with a single rule the

average accuracy benefit of Loc/Time+ (51%) over that of white lists (35%) is marginally significant, whereas with a global limit it took three rules to reach that level. With a two-rule limit the accuracy benefits of Loc/Time+ (54%) and Loc/Time (50%) over that of white lists are significant and marginally significant, respectively. This demonstrates how expressive privacy mechanisms are likely to be more effective under user-burden considerations in more specialized services.

4.4.4 Results related specifically to location sharing with advertisers

As location sharing continues to grow as a social phenomenon (e.g., the recent launch of Facebook’s Places continues to move this trend toward the mainstream), we are already beginning to see location-based coupons being offered to users. Foursquare, a popular mobile location-sharing application, is currently leading this push by allowing small businesses and national chains (e.g., Starbucks) to offer recurring, frequency-based, and loyalty-based coupons to over three million users [99]. Businesses that register with Foursquare are also given access to personally identifiable information about users that visit their locations, such as the names of the most recent and frequent visitors. In this subsection, we delve deeper into our subjects’ attitudes toward sharing location information with advertisers.

The pre-study survey contained one question related to advertisers, asking participants to “rate how comfortable you would be if **advertisers** (e.g., in order to send you promotions or coupons) could view your location,” either always, at user-specified times, or user-specified locations. On a 7-point Likert scale, where 1 was labeled “Not comfortable” and 7 was “Fully comfortable,” users reported an average of 2.6 for always, 3.6 with specified times, and 4.3 with specified locations. Both time and location specifications made users significantly more comfortable ($p < 0.01$ for both, paired t-tests, time: $t = 3.11$, $df = 24$; location: $t = 4.28$, $df = 24$) and location specifications were significantly more comforting than time ($p < 0.01$, $t = 2.98$, $df = 24$).

The post-study survey contained several questions related to advertisers. The first asked “how bad” it would be if a user’s location was disclosed to advertisers when they did not want it to be, and also the reverse (i.e., a non-disclosure when disclosure was wanted). On a 7-point Likert scale, where 1 was labeled “Not bad at all” and 7 was “Very, very bad,”

participants reported an average discomfort level of 1.67 for mistakenly not-disclosing, and an average discomfort of 4.74 for mistakenly disclosing a location. These results suggest that, as expected, a missed opportunity is only a minor concern to our users, whereas disclosing a privacy-sensitive location to advertisers has a significantly higher cost.

We also asked users what the most important factors would be in allowing advertisers access to their locations. The results from this question, again on a 7-point Likert scale from “Not important” to “Very, very important,” are displayed in Figure 4.11. From the reported responses, a user’s location and the quantity of ads received mattered significantly more than the brand of the advertisers and time of day.




Figure 4.11: User responses on qualities of advertisers which would impact their future location-sharing decisions. Answers were reported on a 7-point Likert scale, from 1 (Not important) to 7 (Very, very important). Averages and 95% confidence intervals are shown.

In Figure 4.12 we see that, as with the Facebook group, as the cost of mistakenly revealing a location increases, the policies become more restrictive and the average time shared with advertisers decreases. However, more expressive privacy mechanisms, such as Loc/Time+,




Figure 4.12: The average percentage of time shared with the Advertiser group, under each of the different privacy mechanisms, while varying the cost associated with mistakenly revealing a location from $c = 1$ to 100.

resist this decrease, and allow policies that maximize the amount of time shared while preventing high-cost mistakes.

For even moderate values of c , such as $c \geq 15$, more expressive mechanisms, such as Loc/Time+ and Loc/Time, result in nearly three times as much sharing as Opt-in (i.e., white lists), and this difference is statistically significant ($p < 0.05$ for Loc/Time+ and $p < 0.1$ for Loc/Time). This substantial increase in sharing with large values of c is particularly relevant, given that our subjects reported being very concerned about sharing locations marked private with advertisers in our post-study survey.

Additionally, we find that the increases in sharing from more expressive privacy mechanisms can be realized, even if users are only willing to make a small number of rules. As displayed in Figure 4.13, with $c = 20$, we see a substantial increase in the percentage of time a user would share his or her location with only a single rule. With two rules the differences between the expressive mechanisms, Loc/Time+ and Loc/Time, and Opt-in are statistically significant ($p < 0.05$) and marginally significant ($p < 0.1$), respectively. And, as the cost of mistakes increases, the increase in sharing under more expressive mechanisms with small numbers of rules is even more dramatic. For example, when $c = 100$ we see an almost three




Figure 4.13: The average accuracy (bars indicate 95% confidence intervals) achieved by each of the different privacy mechanisms for the Advertisers group, while varying a limit on the number of rules in a policy that apply to advertisers only. Results are shown for two different values of c , $c = 20$ and $c = 100$.

times increase in sharing over Opt-in with a single Loc/Time+ rule, and this increase is statistically significant ($p < 0.05$).

4.5 Conclusions and future research

Over the past few years we have seen an explosion in the number and different types of applications that allow individuals to exchange personal information and content that they have created. While there is clearly a demand for users to share this information with each other, they are also demanding greater control over the conditions under which their information is shared.

This chapter presented the results from a user study that tracked the locations of 27 subjects over three weeks to collect their stated privacy preferences. Throughout the study, we collected more than 7,500 hours of data. In contrast to some earlier research that identified the requester’s identity [53] and user’s activity [52] as primarily defining privacy preferences for location sharing, we found that there are a number of other critical dimensions in these preferences, including time of day, day of week, and exact location.

We characterize the complexity of our subjects’ preferences by measuring the accuracy of different privacy mechanisms. We considered a variety of mechanisms with differing levels and forms of expressiveness.

As one might expect, we found that more expressive privacy mechanisms, such as those that allow users to specify both locations and times at which they are willing to share, were significantly more accurate under a wide variety of assumptions. More surprising was the magnitude of the improvement — in some cases we found an almost three times increase in average accuracy over that of white lists. These findings were also consistent with our pre-study survey, where subjects reported being significantly more comfortable with the prospect of sharing their location using time- and location-based rules.

We also measured the amount of time that our subjects would have shared their location under each of the different privacy mechanisms. We found that more expressive mechanisms also generally lead to more sharing. This result, which may at first seem counter intuitive, is due to the fact that users generally tend to err on the safe side, and restrict access with

simpler mechanisms. This suggests that offering richer privacy settings may make services more, not less, valuable, by encouraging privacy-sensitive users to share more.

One practical implication of our work is that white lists appear to be very limited in their ability to capture the privacy preferences revealed by our study. This, in combination with the fact that white lists are the only privacy mechanisms offered by most location-sharing applications today (with the notable exception of Locaccino developed by our research group at CMU, which offers all of the expression types we discussed) [122], suggests that the slow adoption of these services may, in part, be attributed to the simplicity of their privacy settings.

Clearly, as privacy settings become more complex, users may have to spend more time specifying their preferences. To address this, we also examined the impact of the different privacy mechanisms under varied assumptions regarding the amount of effort users would be willing to exert while creating their policies. Our findings suggest that, while limiting policies to a small number of rules dampens the accuracy benefits of more expressive mechanisms, they generally remain substantially more accurate than white lists.

The user study presented in this chapter also demonstrates a general methodology for characterizing the tradeoffs between more expressive mechanisms and accuracy (or efficiency) in a number of privacy- and security-related domains. At a high level, the methodology involves i) collecting highly detailed preferences from a particular user population, ii) identifying policies for each subject under a variety of different privacy or security mechanisms, and iii) comparing the accuracy of the resulting policies under a variety of assumptions about the sensitivity of the information and tolerance for user burden.

The findings in this chapter also open several avenues for future work. One avenue involves exploring additional dimensions of privacy preferences. For example, we can study mechanisms that allow users to control the resolution at which location information is provided (e.g., neighborhood, city, or state), or that grant access based on the user's proximity to the requester. We can elicit more detailed information about how bad it would be if each location were to be revealed. We can also investigate the impact of accuracy models that are richer in terms of their tolerance for error. For example, we can use models with costs for mistakenly revealing a location that depend on the subject, the requester, the time of day, or the location in question. With a larger study size, we could also consider how results

vary by demographic subgroup.

We examined the impact of a rule limit on the accuracy of more expressive privacy mechanisms, but we still assumed that users would be able to identify the most accurate possible rules subject to this limit. This opens up another avenue for future work: accounting for additional cognitive limitations, such as bounded rationality [136], to address issues that challenge this assumption. One potential method for accomplishing this would be to study the behavior of real users of a location-sharing application that offers all of the different expression types discussed in this chapter, such as *Locaccino*. In such a study we could provide actual users with different privacy mechanisms and measure the amount of sharing that occurs under each. We could then compare actual user behavior to the predictions of our models, and better characterize the difference between what is predicted by our analysis and what users will actually do in practice.

Another interesting aspect to consider in future work is the value of “negative information.” For example, a user who shares his or her location everywhere other than at home is implicitly sharing it at all times, since a requester can infer from a denied request that the user is at home. In our study, this was not a concern since no actual sharing was being done. However, it would be interesting to consider how such issues affect users’ attitudes towards location sharing within a service that is constantly tracking them.

Finally, there are also legal and policy implications for our work. For example, the information that is protected by the mechanism from other users may not be protected from certain legal entities. In this case, there are two approaches: either the privacy mechanism can be placed on the tracking device itself, which would prevent the information from being recorded, or policies can be enforced that purge the stored data on a regular basis. These factors may influence user preferences and would be interesting to consider in future work.

4.6 Appendix: Survey materials

In this appendix, we provide the exact wording of our pre- and post-study surveys and survey questions. Materials were created using the online tool SurveyMonkey (<http://www.surveymonkey.com>) and have been translated from their screen versions here as accurately as possible.

4.6.1 Pre-study survey

Thank you for your interest in the Location Trails study. To participate, you must be a Carnegie Mellon University affiliate and have an AT&T/Cingular or T-Mobile phone with a SIM card.

If you are selected for this study, you will be asked to use a Nokia N95 smartphone (which we will provide for the duration of the study) as your regular cell phone, and to carry it around with you for 3 weeks (21 days).

At the beginning of the study we will host a brief info session, where we have you come to a room at the Carnegie Mellon University Center, give you your N95 for the duration of the study, load your SIM card into it, and give you a brief demo of our system. We expect this to take approximately 20 minutes.

During the study, you will be required to answer a 10 - 15 min online survey, each day. This survey will ask you some simple questions regarding the places you visited, that day. This can be done by you, online, whenever you see fit.

If you complete the study, you will receive an Amazon.com gift certificate for \$50. This constitutes a payment of \$10 for the first two weeks, a \$20 bonus for the final week and successful return of the phone, and a \$20 reimbursement for one month of an unlimited cellular data plan (if you already have an unlimited data plan you will still receive the full \$50).

1. What is your name?
2. What is your Andrew ID?
3. What is your gender?

4. What is your age?
5. What is your status at Carnegie Mellon?
 - Undergrad
 - Grad. student
 - Staff
 - Faculty
6. Who is your current cell phone service provider?
7. What is the brand and model of your cell phone? (e.g., Motorola Razr, HTC Touch, LG Chocolate.)
8. Are you willing to use a Nokia N95 as your primary cell phone for the duration of the study (3 weeks)?
9. Does your current cell phone plan include data usage (email, Internet)? Having a texting plan is not data.
 - Type of data plan you currently have? (if you know):
10. Would you be willing to purchase an unlimited data plan for one month, for the duration of this study. (For this you will be reimbursed \$20)
11. What kind of computer do you have?
 - Desktop
 - Laptop
12. How often do you use Facebook?
 - Don't have an account
 - Rarely
 - Weekly
 - Daily

- Several times per day

13. Please rate how comfortable you would be if your *close friends and family* could view your location:

	Not comfortable at all				Fully comfortable (Not a problem)		
	1	2	3	4	5	6	7
Anytime							
At times you have specified							
At locations you have specified							

14. Please rate how comfortable you would be if your *anyone you are friends with on Facebook* could view your location:

	Not comfortable at all				Fully comfortable (Not a problem)		
	1	2	3	4	5	6	7
Anytime							
At times you have specified							
At locations you have specified							

15. Please rate how comfortable you would be if your *anyone at CMU* (e.g., anyone in the CMU Facebook network, even people you are not friends with) could view your location:

	Not comfortable at all				Fully comfortable (Not a problem)		
	1	2	3	4	5	6	7
Anytime							
At times you have specified							
At locations you have specified							

16. Please rate how comfortable you would be if your *advertisers* (e.g., in order to send you promotions or coupons) could view your location:

	Not comfortable at all				Fully comfortable (Not a problem)		
	1	2	3	4	5	6	7
Anytime							
At times you have specified							
At locations you have specified							

17. Please rate how much you agree or disagree with the following statements.

	Strongly disagree				Strongly agree		
	1	2	3	4	5	6	7
I can solve most technical problems I am confronted with.							
Technical equipment is often difficult to understand and master.							
I enjoy solving technical problems.							

18. In what instances do you think having people view your location would be useful? Please list at least 1 case.

4.6.2 Post-study survey

Thank you for participating in the Location Trails study. Upon completion of the exit survey, you will receive, via email, a gift certificate for Amazon.com.

1. What is your Andrew ID? (We need this so we can send you your Amazon.com gift certificate!)
2. Do you feel there are any groups of people who might want to see your location other than the groups we asked about (e.g., other than Close Friends & Family, Facebook Friends, Anyone at CMU, and Advertisers)? If yes, please list them?
3. Who did you think of when considering the Advertisers group? In what cases were you willing to share your location with them?

4. Please indicate how important the following factors would be in determining whether or not you would be willing to share your location with advertisers in the future:

	Not important at all				Very, very important		
	1	2	3	4	5	6	7
The time of day							
The location you are at							
The type of advertisers							
The brand of the advertisers							
The type of product being advertised							
The number of ads you receive							

5. In our study your locations were not actually given to the groups we asked about, even if you indicated that you would have been comfortable sharing it. How often do you believe you would have answered the Location Trails questions differently if we had actually been giving out your location to the groups we asked about?
- Never
 - Rarely
 - Some of the time
 - Most of the time
 - All of the time
6. Based on your experiences in this study, please rate how concerned you were, overall, for your privacy when using a location-sharing application.
- 1 - Not concerned, 2, 3, 4, 5, 6, 7 - Extremely concerned?
7. Assuming you were using a real location-sharing system, how bad would it be if the system accidentally __SHARED__ your location with the following groups when you __DID NOT__ want it shared?

Not bad at all

Very, very bad

	1	2	3	4	5	6	7
Close Friends & Family							
Facebook Friends							
Anyone at CMU							
Advertisers							

8. Assuming you were using a real location-sharing system, how bad would it be if the system __DID NOT__ share your location to someone in the following groups when you __WANTED__ it shared?

Not bad at all

Very, very bad

	1	2	3	4	5	6	7
Close Friends & Family							
Facebook Friends							
Anyone at CMU							
Advertisers							

9. For the past few weeks, you were using an online location-sharing application called Locaccino. Please answer the following questions about your experiences with the technology. Please select whether you agree or disagree with the following statements:

Strongly disagree

Strongly agree

	1	2	3	4	5	6	7
It was easy to carry around and use the Nokia N95.							
It was easy to keep the Locaccino client running.							
The locations provided about me were accurate.							
It was easy to answer the survey questions on Location Trails each night.							

Chapter 5

Expressiveness in Price Catalogs

5.1 Introduction

Business-to-customer retail sales account for nearly four trillion dollars in the United States annually, and the percentage of this shopping done online increased three-fold between 2002 and 2007 [145, 146]. Yet, despite the increased computational power, connectivity, and data available today, most online and brick-and-mortar retail mechanisms remain nearly identical to their centuries-old original form: item-only catalog pricing (i.e., take-it-or-leave-it offers). These are the default of B2C trade and are used by massive online retailers like Amazon, Best Buy, and Dell. However, they are fundamentally inexpressive because they typically do not allow sellers to offer discounts on different combinations, or *bundles*, of items.

Recently, some electronic retailers have started offering large numbers of bundle discounts (e.g., motherboards and memory at the popular computer hardware site, New Egg, and songs or albums on music sites), and brick-and-mortar retailers often offer bundle discounts on select items, such as food and drinks. Such discounts make the item-only catalog more expressive, and can be viewed as part of the general trend toward increased expressiveness in economic mechanisms. Increases in expressiveness have been shown to yield better outcomes in the design of general economic mechanisms, as we discussed in Chapter 2 [20, 23], and in a number of specific domains such as sourcing auctions [125], advertisement markets, as we discussed in Chapter 3 [21, 152], and privacy mechanisms, as we discussed in Chapter 4 [19, 85].

Researchers in economics, operations research, and computer science have studied issues surrounding choosing prices and bundles in various types of catalog settings for decades. However, this work has either been i) largely theoretical in nature rather than operational (e.g., [2, 7, 57, 96, 133]), ii) focused on specific types of customer survey data which is not available in many applications (e.g., [69, 81, 120]), or iii) focused on specific sub-problems (e.g., pricing information goods [10, 37, 72, 86, 156], item-only pricing [12, 18], or unit-demand and single-minded customers [67]). (Much of this related work is discussed at more length in Chapter 6.) Despite the ability to collect substantial amounts of data about actual customer responses to different pricing schemes, retailers in most domains are still lacking practical techniques to help them identify promising bundle discounts to offer.

In this chapter, we introduce an automated framework that suggests profit-maximizing prices, bundles, and discounts, the first, to our knowledge, to attempt bundle discounting

using shopping cart data. Our framework uses a pricing algorithm to compute high-profit prices and a fitting algorithm to estimate a customer valuation model. As new purchase data is collected, it is integrated into the model fitting process, leading to an online technique that continually refines prices and discounts.

In Section 5.5, we conduct computational experiments that test each component of our framework individually and one set that tests the framework as a whole. Our results reveal that, in contrast to the products typically suggested by recommender systems, the most profitable products to offer bundle discounts on appear to be those that are only occasionally purchased together and often separately. We also use data from a classic shopping cart generator [5] to estimate the gains in profit and surplus that can be expected by using our framework in a realistic setting. We conservatively estimate that a seller with shopping cart data like that of the generator, who already has optimally priced items, can increase profits by almost 3% and surplus by over 8% using only bundles of size two (even if he has a thousand items for sale). All of our results taken together suggest that this line of work could have material practical implications.

The setting we consider in this chapter involves a seller with m different kinds of items who wishes to choose a set of prices to offer on different combinations of those items to one customer at a time.¹ However, we generalize our framework to consider settings with more than one customer by measuring expectations for profit and revenue, which implies that item prices cannot depend on the identity of the customer. We also consider the special case where a seller can only offer discounts on bundles and must hold the item prices fixed for some exogenous reason (e.g., due to existing policies or competition). We also assume the seller has a cost function that can be approximated by assigning each item a fixed cost per unit sold (in the case of digital goods, which have no marginal cost to produce, we assume the seller can estimate some form of amortized cost), and his goal is to maximize expected *profit* (revenue minus cost). The seller chooses a price catalog, $\pi(b)$, which specifies a take-it-or-leave-it price for each bundle, b , of items. In an item-priced catalog, the price of a bundle is the sum of its parts. (We will be studying richer price catalogs than that, but we still will

¹For settings where the seller has a very large value of m (e.g., supermarkets or large online retailers, which can have hundreds of thousands or millions of different items), we can perform our analysis independently on significantly smaller subsets of the seller's full offering, assuming customers make decisions about the subsets independently.

not be pricing each bundle separately in order to keep the process tractable.) The customer has a valuation, $v(b)$, for each bundle b , and chooses to purchase the bundle that maximizes her *surplus* (valuation minus price). We make the usual assumption of free disposal (i.e., the value of a bundle is at least as much as the value of any sub-bundle).² We measure expected values of revenue, seller’s profit, surplus, and *efficiency* (buyer’s surplus plus seller’s profit).

5.2 Item pricing can be arbitrarily inefficient

The following result provides a theoretical motivation for our work on bundling by demonstrating how our theoretical framework from Chapter 2 can be used to characterize the inefficiency of the item-only catalog.

Proposition 15. *Consider any catalog pricing setting with at least two items, a and b . Using an item-only price catalog, the seller cannot semi-shatter the two pairs of outcomes where the customer buys $\{\{a\}, \{b\}\}$ and $\{\{a, b\}, \emptyset\}$.*

Proof. To adapt our framework to this domain, we construct a mapping from each of the theoretical entities in the catalog pricing setting to an equivalent entity in our framework. Under the mapping, the seller’s type, which is unknown to the customer, determines his or her cost function. The customer’s type, which is unknown to the seller, determines his or her valuation for every combination of items. For a fully expressive catalog, the seller’s expression space includes every possible offering of prices on every combinations of items. For an item-only catalog, the seller’s expression space consists of offerings on individual items only. We assume that the buyer’s expression space is the same as her type space, rather than having the expression space be simply a choice of bundle, and that, given that type, the outcome function chooses the outcome corresponding to a surplus maximizing bundle for the customer. This is a more appropriate mapping than what may seem like a natural

²Contrary to some prior work, which assumes customers have valuations for items only [43], or that valuations are fixed and known in advance [69], here the customer can have valuations over bundles (though in some of our experiments we restrict the complexity of these valuations to ensure tractability for larger values of m), with a valuation for each bundle that is drawn from a probability distribution. (Since the seller may not always know this distribution ahead of time, we also discuss methods for estimating the distribution from historical purchase data in Section 5.4.)

alternative (i.e., mapping the customer's expressions directly to bundle choices) because it ensures that the outcome function is sensitive to the expressions of both the buyer and the seller. For example, under the alternative mapping, holding the buyer's expression fixed and drastically changing the seller's prices would not change the outcome chosen.³ \square

Together with prior results on the role of semi-shattering from Chapter 2, this implies that the item-only catalog is arbitrarily inefficient for some cost functions and valuation distributions, even if the buyer and seller are trying to maximize efficiency.

However, typically in catalog settings, the seller is also the mechanism designer, and will offer more expressive catalogs only if that results in greater expected profit. Also, such a seller will choose prices that maximize profit rather than efficiency. In order to address these issues, we will now demonstrate how we can adapt the computational methodology described in Chapters 3 and 4 to determine the cases when it is in a seller's best interest to use a more expressive price catalog. We will also explore the implications of doing so on the economy as a whole. As a side product of this investigation, we will develop and compare a number of different algorithms for computing high-profit prices for a given valuation distribution, and methods for estimating this valuation distribution from historical purchase data.

5.3 Searching for profit maximizing prices

To study the impact of the item-only price catalog's inexpressiveness in practice, we first develop pricing algorithms that can determine the seller's profit-maximizing prices for a given type of catalog, cost function, and distribution over customer valuations. These algorithms will enable us to measure the expected profit, efficiency, and surplus of different catalog mechanisms for various settings, and allow us to identify characteristics of valuation distributions where the economy is particularly hurt by item-price-only catalogs. The algorithms will also be of practical use when combined with the methods described in the following section for learning customer valuations from historical purchase data.

³Essentially, what we have done is encode the surplus-maximization process in the outcome function. This gives us a more meaningful way to study the mechanism's expressiveness, and is without loss of generality, assuming the customer will always choose the bundle that maximizes his or her surplus.

At the heart of our algorithms is a probability distribution over outcomes for a given price catalog. For any given bundle, b , and price catalog, π , we assume the seller can estimate $P(b|\pi)$, the probability that the customer will buy b . Estimating $P(b|\pi)$ is non-trivial since its domain is exponential in the number of items and it can be fairly complex. For now, we will not assume that the estimated probability function, \hat{P} , has any particular form, other than being a *valid* probability distribution (i.e., the purchase probabilities of all bundles, including the empty bundle, sum to one for every possible catalog). In the next subsection, we will describe one method that we have developed for estimating such a probability function from historical purchase data.

Each of our algorithms takes as input an estimate of the probability function, \hat{P} , the seller's cost function, $c(b)$, a set of priceable bundles, \hat{B} (determined by the type of catalog), lower and upper bounds on the price of each bundle, $L(b)$ and $U(b)$ (also determined by the type of catalog, and can be used to ensure certain prices are fixed), and a seed price catalog, $\pi^{(0)}$ (which need not be intelligently generated). We assume that the algorithm can choose any arbitrary prices for the different bundles as long as the price of a bundle is no greater than the sum of any collection of sub-bundles that contain all of its items.⁴ The algorithms each call \hat{P} repeatedly with different candidate catalogs in order to try to identify the one with the highest expected profit: $\max_{\pi} \sum_{b \in B} \hat{P}(b|\pi) \times (\pi(b) - c(b))$.

- Exhaustive pricing (EX): For each priceable bundle, $\hat{b} \in \hat{B}$, this algorithm discretizes the space between $L(\hat{b})$ and $U(\hat{b})$ into k evenly-spaced prices and checks the expected profit of every possible mapping of prices to priceable bundles. It finds an optimal solution (subject to discretization), but is intractable with more than two items and even with two items if k is too large. For a fully expressive catalog (i.e., one where each bundle is priced separately) with m items, this algorithm calls \hat{P} with k^{2^m-1} different catalogs, and \hat{P} can, itself, be costly to compute. Thus, we propose this algorithm be used primarily as a tool to compare results with the other algorithms on small instances.
- Hill-climbing pricing (HC): Starting with the seed catalog, this algorithm computes the

⁴This essentially ensures the catalog is consistent so that a customer cannot get a better price on a bundle by purchasing its components in some other combinations. This is similar to the free disposal assumption for customer valuations discussed earlier but applied to the seller's catalog.

improvement in expected profit achieved by adding or subtracting a fixed Δ from each priceable bundle, which involves $2|\hat{B}|$ calls to \hat{P} , in each step. It updates the catalog with the change that leads to the greatest improvement, and repeats this process until there are no more improving changes. The resulting catalog is returned, and, since the catalog is only updated when an improvement is possible, it is guaranteed to have the highest observed expected revenue.

- Gradient-ascent pricing (GA): Starting with the seed catalog, this algorithm computes the gradient, or partial derivative, of the expected profit function, which involves $|\hat{B}|$ calls to \hat{P} in each step. The partial derivative, $\vec{d}(b)$, of the expected profit function with respect to a bundle, b , is estimated by measuring the change in expected profit when a fixed Δ is added to $\pi(b)$. The resulting vector of derivatives, \vec{d} , is normalized to sum to one, and the algorithm updates its best candidate catalog by adding $\vec{d}(b) \times \Delta$ to the price of each priceable bundle. The algorithm continues this process until no more improvements in expected profit are possible. The resulting catalog is returned, and, as with the hill-climbing algorithm, it is guaranteed to be the one with the highest expected revenue that was explored throughout the search. In our experiments, this algorithm achieved near-optimal expected revenue on most instances, while performing poorly on a few, with a relatively few number of calls to \hat{P} .
- Pivot-based pricing (PVT): This algorithm generalizes hill-climbing by searching for the best adjustment to the current prices of up to k bundles at a time. For each k -or-less-sized combination of priceable bundles, β , this algorithm measures the change in expected profit from simultaneously adjusting all the prices in β . Each price can be incremented by Δ , decremented by Δ , or not changed. At each step, the algorithm tests all of those possibilities and selects the one that increases expected profit the most. The hill-climbing algorithm above is a special case of this where $k = 1$. However, for larger values of k it generalizes that algorithm to consider more complex types of price adjustments. This process involves $\binom{|\hat{B}|}{k} \times (3^k - 1)$ calls to \hat{P} at each step. With two products and $k = 2$, Table 5.1 illustrates all of the gradients this algorithm would test during each step. Each group of three cells in the table, enclosed in a brackets, represents a gradient, and an arrow indicates the direction of the corresponding bundle's (indicated by the column header) price movement. Even with $k = 2$, our early

a	b	$\{a, b\}$	a	b	$\{a, b\}$	a	b	$\{a, b\}$	a	b	$\{a, b\}$	a	b	$\{a, b\}$
$\langle \uparrow$		\rangle	$\langle \downarrow$		\rangle	$\langle \uparrow \uparrow$		\rangle	$\langle \uparrow$	\uparrow	\rangle	\langle	$\uparrow \uparrow$	\rangle
\langle	\uparrow	\rangle	\langle	\downarrow	\rangle	$\langle \uparrow \downarrow$		\rangle	$\langle \uparrow$	\downarrow	\rangle	\langle	$\uparrow \downarrow$	\rangle
\langle		$\uparrow \rangle$	\langle		$\downarrow \rangle$	$\langle \downarrow \downarrow$		\rangle	$\langle \downarrow$	\downarrow	\rangle	\langle	$\downarrow \downarrow$	\rangle

Table 5.1: An illustration of all of the gradients considered by the pivot algorithm during each step with two products and $k = 2$. Each group of three cells in the table (enclosed in angle brackets) represents a gradient, an arrow indicates the direction of the corresponding bundle’s price movement. A cell with no arrow indicates no movement in the corresponding bundle’s price.

tests show this is the only one of the algorithms (other than the exhaustive one), that achieves optimal expected revenue on nearly every instance we have explored.

5.4 Estimating a rich customer valuation model

The problem of estimating a customer valuation model from historical purchase data is an essential part of our bundling framework because it allows us to use the pricing algorithms presented in the previous section in a practical setting. It is also a problem of interest in its own right, as it extends the classic *market basket analysis* problem first introduced by Agrawal *et al.* in 1993 [4]. Market basket analysis is a commonly studied data mining problem that involves counting the frequencies of different bundles in a collection of customer purchase histories. Simply counting these occurrences can be challenging when there is a large set of items and each customer buys several of them at once. Almost all of the work on this problem has focused on building recommender systems that suggest products frequently purchased together. Many algorithms have been developed for finding bundles with various statistical properties, including one that was developed and patented by Google co-founder Sergey Brin and others [35]. However, as our experiments in Section 5.5 show, our framework predicts that the most profitable items to bundle are those with the opposite profile.

The valuation modeling problem that we consider extends the market basket analysis problem to involve predictions about what would happen to the purchase frequencies under

different price catalogs. (There has been significant recent progress on inferring valuation distributions from bids or other indications of demand in a variety of applications (e.g., [9, 16, 81, 84, 108, 148]), but that work focuses primarily on using bids in auctions or survey information to estimate valuation distributions.)

The inputs to the two problems are essentially the same, although in the case of our valuation problem we include the price catalogs that were on offer at the time of purchase, which can provide additional information about sensitivities to price changes. The close relationship between these two problems allows us to use a classic data generator for the market basket problem in our experiments.

5.4.1 Deriving the maximum likelihood estimate

For the valuation modeling problem, we are given a set of historical purchase observations, $D = \{\langle b_1, \pi_1 \rangle, \langle b_2, \pi_2 \rangle, \dots, \langle b_n, \pi_n \rangle\}$, where each observation, i , includes a bundle that was purchased, b_i , by a distinct customer, i , and the prices of all bundles at the time, π_i . In the case of item-only pricing, a bundle's price is the sum of the prices of the items it contains. (In practice, it is likely that many of the observations will have the same π .) We assume that these purchases are made based on each customer's surplus-maximizing behavior with valuations drawn from an underlying valuation model. We also assume that each purchase is independent of all others since we consider each observation to be from a distinct customer. We will now show that, under these assumptions, the maximum likelihood estimate (i.e., model that maximizes the likelihood of the data) for the customer valuations yields a \hat{P} that matches the observed purchase frequencies as closely as possible. For shorthand, we denote $\hat{P}(B = b_i | \pi_j)$ as \hat{P}_{ij} . The log likelihood of the data given \hat{P} , $\ell(D | \hat{P})$, is then given by the following.

$$\begin{aligned}\hat{P}_{ij} &= \hat{P}(B = b_i | \pi_j) \\ L(D | \hat{P}) &= \prod_i \hat{P}_{ii} \\ \ell(D | \hat{P}) &= \sum_i \log(\hat{P}_{ii})\end{aligned}$$

We can rewrite $\ell(D | \hat{P})$ by aggregating over catalogs and bundles instead of data points. For short, we denote the number of observations containing catalog j as D_j , and the number

of observations containing bundle i and catalog j as D_{ij} .

$$\ell(D | \hat{P}) = \sum_j \left[\sum_{b \in B} D_{bj} \log(\hat{P}_{bj}) + (D_j - \sum_{b'} D_{b'j}) \log(1 - \sum_{b'} \hat{P}_{b'j}) \right]$$

Next, we take the partial derivative of ℓ with respect to a given value of \hat{P} , set it equal to zero, and solve for the point where the data likelihood is not changing:

$$\begin{aligned} \frac{\partial \ell}{\partial \hat{P}_{ij}} &= \frac{D_{ij}}{\hat{P}_{ij}} - \frac{D_j - \sum_{i'} D_{i'j}}{1 - \sum_{i'} \hat{P}_{i'j}} \\ 0 &= \frac{D_{ij}}{\hat{P}_{ij}^*} - \frac{D_j - \sum_{i'} D_{i'j}}{1 - \sum_{i'} \hat{P}_{i'j}^*} \\ \frac{D_{ij}}{D_j} (1 - \sum_{i'} \hat{P}_{i'j}^*) &= \hat{P}_{ij}^* (1 - \sum_{i'} \frac{D_{i'j}}{D_j}) \end{aligned}$$

If we assume all the bundle probabilities other than i are equal to the percentage of the data in which they appear under each catalog (i.e., $\forall i' \neq i, \hat{P}_{i'j}^* = \frac{D_{i'j}}{D_j}$), then $\frac{D_{ij}}{D_j}$ is the unique solution for \hat{P}_{ij}^* .

5.4.2 Fitting the valuation model to purchase data

The valuation model we will fit allows for normally distributed valuations on each item, pair-wise covariance between valuations for items, as well as normally distributed terms for complementarity (or substitutability in case such a term is negative). This model significantly generalizes prior ones [39, 43, 133, 150] by allowing for heterogeneous complementarity and substitutability between products.

Specifically, our model parameters include a mean and variance for each priceable bundle in \hat{B} and covariances between individual items' valuations. While the draw, $x_{\{i\}}$, from the distribution of an item i represents that item's valuation, $v(\{i\})$, to the customer, a draw from the distribution for a bundle b of two or more items represents a complementarity bonus (or substitutability penalty if negative). The valuation for a bundle is then the sum of the draws of all the bundles (including individual items) it contains: $v(b) = \sum_{b' \subseteq b} x_{b'}$. Under this

model, a customer’s valuation can be thought of as a hyper-graph where each (hyper-)edge is associated with a real-valued random variable representing the valuation bonus or penalty for receiving a bundle containing the items connected by the (hyper-)edge. This allows us to model any possible distribution over valuations (without loss of generality),⁵ and can be viewed as a probabilistic generalization of the classic k -wise valuation model introduced by Conitzer and Sandholm for combinatorial auctions [51].

To go from a valuation model to the probability function, \hat{P} , we use a Monte-Carlo method to sample customers (10,000 in our experiments) according to the valuation distribution, and, for a given catalog, we simulate their surplus-maximizing purchasing behavior (taking into account that disposal is free). This simulation is relatively straightforward since items that are not connected by a complementarity or substitutability edge can be considered independently.

In order to identify the model parameters that maximize the likelihood of the observed data, we use a hybrid search technique. It begins by performing a tree search over the variance and covariance parameters. A range for each of these parameters is given as input that is discretized into a specified number of values (in our experiments we use six values per parameter). At each leaf node, a local search is performed to find the means that maximize the data likelihood given the values of the variance parameters at that leaf. In our experiments, we use a pivot-based search, as described in Section 5.3, for this step. The parameter settings resulting in the highest overall likelihood are returned, and in the case of a tie an even mixture of all the tied models is used (i.e., simulated customers are sampled from each with equal probability). Figure 5.1 illustrates this process for two items, a and b .

We found empirically that this technique of first choosing standard deviations using tree search and finding means using local search provided better results than exclusively using tree search or local search for all parameters. This is primarily due to the tight relationship between the appropriate means and standard deviations. Once the standard deviations have been fixed, the best means are relatively easy to identify using local search. However, the best means change drastically with a relatively small change in standard deviations. Using tree search exclusively would also produce good results, but the complexity of such a search

⁵For example, consider a setting with three items where a customer receives a complementarity bonus from any single pair but no additional bonus for more than one pair. Here, we would use complementarity edges between all pairs and a substitutability three-edge connecting all three items to avoid double counting.

makes it infeasible to conduct on fine-grained parameter values.

Most existing shopping cart data involve only a single catalog. They do not include information about customers’ surplus-maximizing behavior under alternative prices, and, thus, are under-specified for the purposes of inferring a valuation model. To address this, on such instances we utilize the existing item prices as an additional piece of information to fit our model. Specifically, among models that fit the observed purchase data (approximately), we prefer models whose profit under the optimal item-pricing for that model is close to the profit of the existing item prices under the model.⁶ Our algorithm does this test once at every leaf of the search tree (after the best model for the leaf has been computed as described above). If there are still several leaves that are (approximately) as good at explaining the purchase data and the existing prices, we use an even mixture over those models.⁷

5.5 Empirical results

We will now discuss the results from several sets of computational experiments that test our pricing and fitting algorithms and reveal some interesting economic insights that emerge as a consequence of our customer valuation model. The next two subsections focus on pricing and fitting two-item instances. The third set of results provides an estimate of the potential achievable by offering bundle discounts on pairs of items from a seller with a thousand items and realistic shopping cart data.

5.5.1 Results with pricing algorithms

The first set of experiments involves using the search techniques described in Section 5.3 to find high-profit prices on a generic class of instances similar to the models used in prior work [39, 43, 133, 150]. We compare the results and performance of the pricing algorithms on symmetric two-item instances where the customer’s valuation for each item is drawn from a normal distribution with mean 0.5 and standard deviation 0.5. We vary the pairwise

⁶One could also compare based on the item-price vector itself, but we prefer the profit-based comparison because it better measures the quality of the original pricing, and we found it to be more stable.

⁷In our experiments we use at most the top five models and fewer if less than five meet the threshold for “approximately as good”.




Figure 5.1: Illustration of the search technique we use for estimating a customer valuation model from historical purchase data. We use a tree search over variances and a pivot-based search over means. Leaves are evaluated based on how closely the corresponding model predicts the observed data and (optionally) how closely the model’s optimal profit matches the profit achieved by existing prices.

Algorithm	E[prof.]	E[eff.]	E[surp.]	\hat{P} calls
PVT	99.99%	97.88%	86.92%	267.25
EX ($k = 15$)	99.24%	97.70%	88.28%	3375.00
GA	97.57%	99.14%	90.78%	49.61
HC	89.76%	93.73%	96.32%	4.08

Table 5.2: The average fraction of the highest expected profit, efficiency, and surplus, as well as the average number of calls to \hat{P} for each of the pricing algorithms described in Section 5.3. For these results, the algorithms were run on symmetric two-item instances.

covariance from -0.25 to 0.25 and we vary the mean of the pair-wise complementarity (or substitutability when negative) term from -1.5 to 0.5 (the standard deviation for this term is held constant at 0.5). Each algorithm (other than the exhaustive one) uses an item-only catalog with all prices set to 0.5 as a seed and a step size $\Delta = 0.05$ to price fully expressive catalogs.⁸ The EX algorithm considers $k = 15$ different prices for each bundle and finds the optimal prices subject to this discretization. The PVT algorithm considers all possible gradients for two item instances.

The following tables and figures illustrate several characteristics of the solutions and performance of the different algorithms for pricing a fully expressive catalog, as well as a variant that only allows for bundle discounts to be offered on the optimal item-only prices (rather than also allowing for changes in item prices).

Table 5.2 reports each algorithm’s average fraction of the highest expected profit, efficiency, and surplus, as well as the average number of calls to \hat{P} over five instances for each parameter setting. The best value in each column is in bold. Other than the unscalable exhaustive algorithm, the pivot-based algorithm is the only one to achieve optimal profit on every instance. Therefore, it is the algorithm we use in the rest of the chapter for pricing. (Gradient ascent also performed well and may scale better for larger instances.)

Figures 5.2 and 5.3 show the increase in expected profit and surplus from allowing sellers

⁸It is possible to improve the performance of all the algorithms, other than the exhaustive one, by starting with a larger step size and repeatedly decreasing it whenever further improvements are impossible at the current size. However, we report results on all algorithms without this improvement for a more meaningful comparison of their performance.

to offer profit-maximizing bundle discounts, while varying the levels of covariance, complementarity, and substitutability. (The values represent averages over five runs but deviate very little.)


For the first set of results, shown in Figure 5.2, we assume the seller holds the item prices fixed at the optimal item-only catalog values to isolate the impact of offering bundle discounts from the potential confound of our system improving the item prices as well. We believe this also represents a practical constraint in many markets and is a policy that sellers are likely to take when first adopting the bundle discounts suggested by our framework. This has the effect of depressing the seller's expected profit gain, but it ensures that the customer surplus cannot decrease.

For the scenarios we consider, the seller's greatest predicted increase in expected profit (about 4.6%) occurs when valuations are highly negatively correlated and the items are slightly substitutable. However, too much substitutability diminishes the predicted profit benefits. Others have also identified negative correlation and substitutability as motivators for offering bundle discounts [39, 43, 150], but they did not use a rich enough valuation model to fully explore the impact of heterogeneous complementarity or substitutability. (That work also did not address the model fitting problem that must be solved to operationalize this insight.)

Unsurprisingly, due to the discount-only pricing we imposed, our results also show a large predicted increase in surplus (averaging around 9%) throughout the parameter space. Together with the seller's predicted increase in profit, this leads to substantial efficiency increases.


Another set of experiments (shown in Figure 5.3) demonstrates that when our system is also free to adjust the prices of the items, additional increases in profit are possible but usually at the expense of the customer surplus.⁹ This may be desirable for the seller in the short term, but maintaining surplus can be an important long-term goal if there are competing sellers.

⁹In some cases, the customer surplus actually decreases by up to 10%, but all values less than or equal to 0% are shown as white dots on the chart for consistency with Figure 5.2.



Results allowing bundle discounting only.

Figure 5.2: The intensity of each dot is the increase in expected profit or surplus achieved by profit-maximizing bundle discounts for different levels of covariance (x-axis) and complementarity (or substitutability) (y-axis), ranging from 0% to 10%. Here, we assume the seller holds the item prices fixed at the optimal item-only catalog values to isolate the impact of bundle discounts.



Results allowing bundle discounting and repricing of items.

Figure 5.3: The intensity of each dot is the increase in expected profit or surplus achieved by profit-maximizing bundle discounts for different levels of covariance (x-axis) and complementarity (or substitutability) (y-axis). Here, we assume the seller has the ability to reprice items as well as offer discounts on bundles. White dots on the surplus graph in this figure indicate that the surplus after repricing was the same or worse than the surplus under the optimal item-only catalog (this is done for consistency with Figure 5.2).

5.5.2 Results with the fitting algorithm

We now present experiments that use the fitting algorithm from Section 5.4 to find models that predict an observed set of purchase data. We allow the search algorithm to consider standard deviations between 0.5 to 3.5 at intervals of 0.5, and we focus on symmetric two-item instances where both items occur with the same frequency in the shopping cart data. (Results on asymmetric instances were similar.)¹⁰

These experiments test our fitting algorithm for the ubiquitous scenario where shopping cart data is accompanied by a single item-only price catalog. As discussed earlier, for these instances we assume the seller’s existing item-only prices are set optimally.¹¹ This provides us with a model that is consistent with both the observed data and the existing pricing.

Figure 5.4 shows the predicted increases in expected profit and surplus achievable by a bundle discount, assuming that the individual items are optimally priced and that at those prices they have the same profit margin (the value of the profit margin does not matter). As in Figure 5.2, we assume the seller can only offer a discount on the existing item prices and cannot change them. When we relax this assumption, we find additional opportunities to increase profit at the expense of customer surplus, consistent with the results shown in Figure 5.3. We consider instances where the item frequencies range from 2.5% to 40% and the *co-occurrence* percentages from 2.5% to 87.5%. We define co-occurrence as the fraction of baskets containing the less frequent item (for symmetric items either can be used) that also contain the other. We also increased our sampling frequency in an interesting area of the parameter space where item frequency is less than 15% and co-occurrence is less than 20%. This is illustrated on each chart by a higher concentration of small points in the bottom left corner. (Again, the values are averaged over five runs but tend to deviate very little.)

¹⁰We also performed a sanity check that our algorithm could fit known models with relatively few sample catalogs. We tested it with 10 random valuation models. Half of them used variances matching the algorithm’s discretized values. Test catalogs were randomly generated and used to compute the algorithm’s average prediction error. In all cases, the fitting algorithm was able to predict purchase frequencies for unseen catalogs with a high degree of accuracy (always less than 5% average prediction error, often less than 1%), using between 2 and 6 samples.

¹¹We also fixed the variance of the complementarity (or substitutability) term to be equal to the average variance of the item valuations since we found that this did not materially affect our results for these instances, and allowed us to run on more instances due to enhanced speed.




Figure 5.4: The intensity of each dot represents the predicted increase in expected profit or surplus achieved by profit-maximizing bundle discounts on single-catalog instances with varying item frequencies (x-axis) and co-occurrence percentages (y-axis), ranging from 0% to 10%. As in Figure 5.2, we assume the seller holds the item prices fixed at the optimal item-only catalog values to isolate the impact of offering bundle discounts from the potential confound of our system improving the item pricing as well.

These results are consistent with those in Figure 5.2, since the seller’s greatest predicted increase in expected profit (about 4.6%) occurs when products are occasionally bought together (co-occurrence probability less than 20%) and frequently bought separately. This set of results also predicts large increases in surplus throughout the parameter space (averaging about 9%), as seen in Figure 5.2. Thus, these results also show a large increase in efficiency.

Taken together, our results illustrate why new techniques are needed beyond those used for building recommender systems, which typically identify items that are commonly purchased or consumed together. In contrast, when it comes to items that can be profitably bundled together at a discount, our framework suggests those with the opposite profile. Our results also explain why recommender systems are highly popular among users: a recommendation can be viewed as a small discount (in the form of time saved), and our framework predicts that even a small discount on highly co-occurring products leads to a substantial increase in surplus.

5.5.3 Results with a shopping cart generator

Our final set of experiments estimates the potential increase in expected profit and surplus achievable by bundling products from a seller with shopping cart data like that generated by Agrawal and Srikant’s classic generator [5]. We use the standard parameters in the generator: for each instance, we generate 10,000 shopping carts with 100-1,000 items (N), 100-2,000 potentially popular bundles (L) of size 2-4 (I), and an average of 2-20 purchases per customer (B). We assume the seller had optimally priced the individual items, and that those prices involved a uniform profit across all items.

Pricing all—or a huge number of—bundles is undesirable for several reasons: i) presenting complex catalogs to customers may be infeasible and/or it may confuse/burden them, ii) it is intractable in terms of computation and information, and iii) even non-overlapping bundles can interact: as one bundle is discounted, some customers might shift from buying other things to that bundle. Therefore, we only consider discounting bundles of two items, and further narrow them down as follows. We only consider item pairs priceable if the items are not directly or indirectly *related* to any other items. We consider two items related if their joint purchase frequency is more than a fixed threshold different than the product of their individual purchase frequencies (we use a threshold of 1% for these experiments). We

N	B	L	I	E[prof.] inc.	E[surp.] inc.
1000	20	2000	4	2.80%	8.34%
1000	20	1000	4	1.10%	3.01%
100	2	200	2	0.89%	2.65%
100	2	100	2	0.15%	0.86%

Table 5.3: The total profit and surplus increases for various parameter settings of a classic market basket analysis generator [5] (values are averaged over five instances for each parameter setting). Here, we assume the seller had priced all items optimally, and at those prices each item was sold for a uniform profit. We also consider discounts only on pairs of items that are unrelated to any other items.

construct a graph where the items are nodes and edges connect items that are related. Then, only connected components of size two and pairs of isolated items are considered priceable.

The profit and surplus increases for each priceable pair are then estimated using the results behind Figure 5.4 and a set of similar results on asymmetric instances. The increase for a given pair is estimated as the average value for the five most similar instances (based on the frequencies of the two items and the bundle). Priceable pairs are then greedily selected to actually be discounted based on their predicted profit increase. Once a pair is selected, all other pairs containing either of the selected items are removed from consideration. Table 5.3 shows the total predicted profit and surplus increases for various parameter settings of the generator (values are averaged over five instances for each parameter setting).

For the standard parameter settings, the first row shows almost 3% profit increase using our algorithms to select pairs of items to bundle and discount. This increase in profit is accompanied by more than an 8% increase in customer surplus, and, thus, a significant efficiency increase. The table also shows that increasing the number of items and potentially popular bundles increases the benefits from our approach. This is because it leads to a sparser relatedness graph and, thus, increases the number of safely priceable items for our algorithms.

These improvement numbers are conservative because they assume that the seller had already priced the individual items optimally. Furthermore, additional improvements may be achievable by using a less conservative method for pricing bundles than our method, which

only prices pairs of items where neither item is related to any other, as discussed above.

5.6 Conclusions and future research

In this chapter, we introduced a framework for automatically mining purchase data and suggesting profit-maximizing prices, bundles, and discounts. It uses a pricing algorithm to compute high-profit prices on items and some bundles, and a fitting algorithm to estimate a customer valuation model. New purchase data can be integrated into the model fitting in an online process that continually refines prices and bundle discounts.

We began by providing a theoretical motivation for bundle discounting based on the theoretical framework developed in Chapter 2: a catalog that only prices items can be arbitrarily inefficient. We then described search algorithms that compute high-profit prices for a given customer valuation distribution. Since it is unlikely that most sellers have an estimate of such a distribution in practice, we introduced a hybrid search technique that uses purchase data to estimate a customer valuation model. Our fitting and pricing algorithms allow us to use a richer valuation model than prior work: in addition to means, variances, and covariances on items, we capture means and variances on complementarity (substitutability).

We reported on computational experiments that examined each component of our framework and finally the complete framework. In contrast to the suggestions of recommender systems, the most profitable products to offer bundle discounts on appear to be those that are only occasionally purchased together and often separately! On realistic shopping cart data, by discounting selected bundles of two items we conservatively estimate almost 3% profit lift simultaneously with an 8% lift in customer surplus. Thus, automated bundle discounting could have significant practical implications.

Some obvious directions for future research include less conservative methods for selecting pricable bundles, discounting bundles of more than two items, and live experiments where the catalogs that we offer serve as demand queries about the customers' valuations that are then incorporated back into our model. These experiments could be carried out similarly to the ones described by Jedidi *et al.* [81], but would involve actual purchases by subjects rather than survey data.

There are also several assumptions made here that could be relaxed in future work. For

example, we assumed that each purchase in the shopping cart data was independent, but it may be possible to develop a model that captures repeat purchases by the same customers. It may also be possible to improve on our results here by offering personalized discounts to take advantage of such repeated purchases. We also assumed that the cost of selling an item could be described by a marginal unit cost. It would be interesting to extend our work here to include considerations of non-linear cost functions (e.g., with large start-up costs) or limited-inventory items. Finally, we assumed that the true customer valuations were drawn from distributions that could be accurately fit by our valuation model. However, it would be interesting to consider the effects of mis-representing these valuations because, for example, they are drawn from a different kind of distribution than the one we use (e.g., log normal instead of normal). It may be that certain pricing algorithms or mechanisms are more robust and better suited to handle such modelling errors.

Chapter 6

Related work

There has been relatively little work on expressiveness specifically. We have discussed several related papers throughout this dissertation. Here we will briefly summarize work on the most closely related topics. This work started in economics and has more recently been studied in computer science.

6.1 Informational complexity

At a high level, related questions go back at least to the 1940s when Hayek argued that in distributed resource allocation, it is not practical to communicate all the distributed information to a central decision maker [70]. In the 1970s, Mount and Reiter [101] and Hurwicz [76, 77] formalized this in their theory of *informational complexity*, which asked the question: at a minimum, how much information must a mechanism’s message space be able to carry in order to accomplish some design goal (cf. [79])? That work focused primarily on the number of real-valued dimensions that were needed. It was well known that in general, as our Proposition 1 shows, the number is always one. To get around Cantor’s theorem that begets Proposition 1, the economists made some technical assumptions (such as local thread- edness [101] or Lipschitz continuity [78]) that precluded a general mapping between \mathbb{R}^n and \mathbb{R}^m . Under these assumptions, Proposition 1 does not apply, and the economists proceeded to compare the informational requirements in different economic settings by comparing the number of dimensions in each agent’s expression. In contrast, our work does not rely on such assumptions. In fact, one of our key points is that the dimensionality of the message space is not the essence of expressiveness. Rather, the essence is how the mechanism is wired to use the different inputs.

6.2 Work based on finding or characterizing equilibria

Another thread of related work has tried to characterize the equilibrium behavior in inexpressive mechanisms in specific settings. The challenge here is that determining equilibrium behavior is usually prohibitively difficult even for the simplest non-trivial mechanisms. Furthermore, when a particular equilibrium is found to have certain properties, one often cannot rule out the possibility of additional equilibria that do not share those properties.

For example, Rosenthal and Wang [119] examined an auction setting where a series of globally interested (with nonlinear preferences over different items) and locally interested bidders (with linear preferences for different items) participate in a set of simultaneous first-price sealed-bid auctions where each auction is for a single item. Taken together, the auctions constitute an inexpressive mechanism. The authors were able to construct an equilibrium for each of two regions of the space of parameter values for the bidder type distributions in their model. They found that these equilibria were inefficient for most of their model parameter space. However, they were not able to rule out the possibility that other equilibria exist (although they have not found any) and they were unable to construct equilibria for some parameter values of their model.

Another example is work by Szentes and Rosenthal [139], who characterized simple efficient equilibria in large inexpressive mechanisms when bidders are identical and each wants to win a specified fraction (more than a half) of the items. The simplicity of this domain illustrates the difficulty in finding equilibria in inexpressive mechanisms. Problems must typically be severely simplified in order to gain traction with analytical or computational techniques.

As further illustration of the difficulty of equilibrium finding, Wilenius and Anderson [155] described a heuristic method for computing approximate equilibrium strategies in first-price sealed bid CAs when bidders either bid on all combinations of items, or on one specific combination and the remaining items individually. They demonstrated the difficulty in finding equilibrium strategies for CAs when they are not dominant-strategy implementable.

All of the work discussed here suggests that there is little hope for a clear general characterization of equilibrium strategies in inexpressive mechanisms.

6.3 Expressiveness in dominant-strategy mechanisms

There has also been some research related to expressiveness issues in dominant-strategy mechanisms. For example, Blumrosen and Feldman [28] studied the problem of designing a dominant-strategy mechanism with a limited number of discrete actions. They showed a trade off between the efficiency of the best possible dominant-strategy mechanism and the number of discrete actions available to the designer. Similarly, Ronen [118] described

methods for achieving near efficiency with limited bidding languages in dominant strategies.

Holzman *et al.* [74] studied CAs where bidders can only bid on restricted sets of bundles. (This is the restricted outcome setting mentioned in Section 2.3.) Their work shows that truthful bidding is a dominant strategy if and only if the restricted bundle set that agents can bid on forms a quasi-field (and VCG payments are used). They defined a worst-case measure of the economic inefficiency that may result from restricting bids to smaller and smaller quasi-fields. Parkes [113] and Nisan and Segal [105] showed that in order to implement VCG payments, a mechanism must elicit enough information to verify the corresponding universal competitive equilibrium prices.

The restriction to studying dominant-strategy mechanisms imposes severe limitations on the types of questions about expressiveness that can be addressed. In particular, uncertainty about others' private information becomes an issue only when considering mechanisms that do not have dominant strategies. As we showed, the larger the possible type space of others, the more expressiveness an agent may need for efficiency. Our results apply to settings where agents do not have dominant strategies (and to settings where they do). Also, our results are not specific to any application, such as a CA.

6.4 Applications of expressiveness in mechanisms

One of the first applications to benefit from expressiveness was strategic sourcing. Sandholm [125, 130] described how building more expressive mechanisms—that generalize both CAs and multi-attribute auctions—for supply chains has saved billions of dollars that would have been lost due to inefficiency. Success with expressive auctions in sourcing has also been reported by others [55, 73, 97]. Schoenherr and Marbert [134] discussed the difficulty faced by business-to-business auction participants in choosing bundles to put up for auction ahead of time. This is a problem that exists because these mechanisms are typically inexpressive: they allow bids on predetermined lots only. If a CA were used instead, the sellers would not have to choose bundles *a priori*: the mechanism would determine the bundles based on the (expressive) bids.

Some work on expressiveness has begun to appear in the context of search keyword auctions (aka sponsored search). Even-Dar, Kearns and Wortman examined an extension of

sponsored search auctions, whereby bidders can purchase keywords associated with specific contexts [60]. Under certain probabilistic assumptions they are able to prove that the system becomes more efficient when this extra level of expressiveness is allowed. We have shown that increasing the expressiveness of today's sponsored search auctions very slightly by allowing for a premium bid for premium slots removes most of the inefficiency of today's design [21]. Also, highly expressive mechanisms have been designed for trading entire advertising campaigns [33, 152]. Milgrom explores the equilibria of sponsored search auctions with limited expressive power (specifically, where bidders submit a single bid to indicate how much they will pay for an ad spot regardless of where it appears on the page) [98]. He finds that by *limiting* expressiveness the auction excludes some bad equilibria. This raises an important counterpoint to our work. We hope that our framework will help us better understand the circumstances under which expressiveness actually helps and when it does not. In another recent paper on sponsored search auctions, Abrams *et al.* studied the impact of inexpressive bids on efficiency [1]. They found that in a specific auction mechanism, inexpressiveness can lead to an arbitrary amount of inefficiency when all bidders are assumed to play the same pure strategy (regardless of what the strategy is). They proceed to show that the same inexpressive mechanism has an efficient *full information* Nash equilibrium even when bidder valuations are more complex. They consider this surprising, but it is consistent with our general result that very little expressiveness is needed for efficiency when agents have no uncertainty (Proposition 7).

Another application area that has received recent attention with regard to expressiveness is wireless spectrum trading. For example, Gandhi *et al.* [62] described a prototype wireless spectrum market mechanism. They stressed the importance of allowing spectrum bidders enough expressiveness to communicate their needs, and demonstrated—using synthetic demand distributions and various *ad hoc* bidder behavior models—that their mechanism has good efficiency properties.

6.5 Related work on bundle pricing and CAs

The first mention of being able to increase revenue via bundling is attributed to economist George J. Stigler in his 1963 discussion of anti-trust Supreme Court rulings over price discrimination via bundling [138] (the issue was whether or not a studio that produced the

films *Gone with the Wind* and *Getting Gertie's Garter* could force theaters to buy the rights to show them together). Bundle pricing in economics has often focused on analyzing two-product settings to provide insight into the way monopolies can improve profits by offering goods in bundles [2, 57, 66, 96, 133]. (One exception is that Armstrong examined n -product settings, but placed severe restrictions on buyers' utility functions [7].) This work provided sufficient conditions on when bundling is profitable and optimal pricing strategies under various assumptions. However, it did not provide generalized algorithms for determining how to price the bundles. Nor did it typically answer the question of how the increase in expressiveness affects the buyers utility or the efficiency of the market as a whole. There have also been some behavioral economics experiments that explored how people actually perceive savings in bundles [157].

Some work on bundle pricing has been done from an operations research perspective as well. For example, Hason and Martin [69] presented a mixed integer program for optimizing bundle prices for a handful of market segments. They assumed that each of the segments can be described by a single value for each bundle, and that the value of every bundle for every market segment is known in advance. They also did not describe how their bundle pricing strategy compared to using item prices. Rusmevichientong *et al.* investigated the problem of pricing different car configurations based on customer survey data collected by GM's Auto Choice Advisor web site [120].

An extensive revenue management literature also exists, but the work in that field tends to focus on pricing individual items in the face of stochastic demand and limited supply. For example, in their seminal book, *The Theory and Practice of Revenue Management* [140], Talluri and Van Ryzin mention bundling as a consideration in revenue management, but discuss dynamic pricing methods for items only (even in the multi-item revenue management settings that were introduced by Gallego and Van Ryzin in 1997 [61]). Two notable exceptions to this focus on item-only pricing are the works of Bulut *et al.* [39] and Venkatesh and Kamakura [150]. These papers consider the problem of selling two products under different bundling policies using customer valuation models similar to those we used in Chapter 5. However, these papers assume customers have a uniform complementarity (or substitutability) term for the items (i.e., they do not model heterogeneity in this term across customers by allowing for variance in the draw for the bundle valuation). Furthermore, Venkatesh and Kamakura make the assumption that item valuations are uniformly distributed, and Bulut

et al. focus on limited-quantity (perishable) products using an additional Poisson process to model customer arrivals. These papers also do not consider the scenario where a seller can only offer bundle discounts on optimal item-only catalogs, they either allow for the items to be repriced [39, 150] or assume item prices are fixed at arbitrary values [39], and they do not provide algorithms for scaling their methods beyond two items. Despite these differences, this work echoes our findings that substitutable and negatively correlated items benefit from bundle pricing, but it does not address the fitting of a valuation model and, thus, does not provide a practical methodology for automatically identifying promising products to bundle in a large catalog.

Another closely related paper in this area is by Jedidi *et al.* [81]. These authors attempt to fit a customer valuation model for two items that is similar to ours using survey data regarding which items subjects would buy under different randomly chosen price catalogs. The work models normally distributed valuations for the items and the bundle, but also includes a parameter for each individual subject in the sample (this is feasible because they consider a setting with less than 100 subjects who each answer several different demand queries). Due to the complexity of this model, the authors shy away from directly estimating its parameters (as we do for our model in Section 5.4), stating that the task is too “difficult.” Instead, they sample thousands of parameter values according to their likelihoods given the survey responses and a hand-crafted prior distribution. The large number of parameters in the model, the need for properly formed priors for all model parameters, and the need for non-trivial amounts of survey data prevent this methodology from being fully automated and scalable.

There have also been several pieces of work specifically on pricing bundles of information goods, where it is usually assumed that customers care only about how many goods are bundled together (i.e., their valuation for a bundle depends only on its size, not its contents) and there are no marginal costs. For example, Kephart *et al.* [86] and Brooks and Durfee [37] described online approaches to pricing in this domain. Additionally, Bakos and Brynjolfsson provided an analytical treatment of this problem with some valuable insights about when bundling is profitable [10]. The operations research literature has also addressed the problem of bundle pricing for information goods. For example, Hitt and Chen [72] and Wu *et al.* [156] consider a bundle pricing mechanism for information goods that allows customers to choose up to M items from a larger pool of N items.

Finally, computer science work on pricing has focused primarily on pricing items rather than bundles, and for “single-minded” customers that desire only one bundle. For example, Balcan and Blum [11] provided online and approximate algorithms for this setting, and Guruswami *et al.* [67] showed that finding the optimal pricing is \mathcal{APX} -Hard. Some work from this community, such as the work by Aggarwal *et al.* [3], considered a more restrictive class of pricing problems called MAX-BUYING, where customers buy the most expensive goods they can afford. Such restricted classes have been shown to be solvable in polynomial time.

Related to bundle pricing, there has recently also been significant work on designing high-revenue CAs (e.g., [49, 82, 109, 131]). Designing for revenue turns out to be much more difficult than designing for efficiency. There have also been recent papers on strategic behavior in CAs, suggesting that increasing expressiveness does not always increase efficiency for some payment functions [41, 42, 93, 106]. (These do not contradict our Theorem 1, since that result allows for a redesign of the mechanism to the allowed level of expressiveness.)

6.6 Related work on location sharing

Location-sharing services are an area of significant growth as consumers gain access to ever cheaper and “smarter” mobile phones. With expanding market share, these services are anticipated to capture a significant portion of the billions of dollars in marketing revenue from the broader class of location-enabled applications [64]. Yet, despite analyst predictions and the growing number of location-sharing applications that have been developed, no service has captured a significant market share.

While high-profile services that are built around location sharing, like Loopt¹ and Google’s Latitude,² seem to dominate the press, neither has been crowned a “killer app.” Dozens of other offerings exist, many built around technology platforms that have allowed easier creation of these applications, including the iPhone SDK,³ and Google’s Android SDK,⁴ as well

¹Loopt. <http://loopt.com/>

²Latitude. <http://www.google.com/latitude>

³iPhone Dev Center. <http://developer.apple.com/iphone/>

⁴Android. <http://code.google.com/android/>

as Yahoo’s FireEagle Platform, which as of March 2010, has 79 applications in its gallery.⁵ The FireEagle platform facilitates privacy-enhanced sharing by allowing users to specify a policy for each service that he or she provides with access. FireEagle allows, just as Google’s Latitude, exact-location or city-level granularity sharing with white-listed entities. However, Tsai *et al.* found that privacy protection through the abstraction of location is rare. Of 89 sharing services surveyed in that work, only 11 provided any control over the granularity of the location disclosure, while over half of the services (50) used white-listing (or, equivalently, black-listing) to protect a user’s location [143]. They also found that more complex privacy-setting types were nearly nonexistent in the landscape at the time, with only 11 services providing group designations, and only two having approvals with expirations. One notable exception was Locaccino,⁶ which was developed by our research group at CMU and allows users to specify time- and location-based rules (these are richer privacy settings than those offered by any commercial service).

Many research groups have developed location-based services, including PARC’s Active Badges [154], ActiveCampus [14], MyCampus [121], Intel’s PlaceLab [71], and MIT’s iFind [75]. However, the research done with these systems rarely reached the point of studying privacy preferences. Instead this work was typically hampered by adoption and technological issues. Work on a Semantic Web framework to capture rich privacy preferences in different context-aware applications, including location sharing applications, was also conducted in the context of CMU’s MyCampus project [121]. This work later led to the development of several other location sharing applications at CMU, including PeopleFinder [122], and most recently Locaccino.

As far back as 2003, users of a diary study cited some concerns about location privacy, stating a preference to not have their phones tracked [14]. A study using the experience sampling method in 2005 found that location-privacy preferences were complex, and “participants want to disclose what they think would be useful to the requester or deny the request” [53]. These findings provide evidence that without more complex privacy-setting types, users will simply shutdown, and deny requests if they cannot specify policies that would lead to useful sharing. One drawback of this research is that much of it focused on laboratory experiments [54, 115] and small group testing [13, 80, 137], where there are

⁵Fire Eagle. <http://fireeagle.yahoo.net/>

⁶Locaccino. <http://locaccino.org/>

minimal privacy concerns given the small number of (often simulated) requests.

As far as we know, there have been only two other field studies that revealed complexity in people's location-privacy preferences. The first, by Tsai *et al.*, found that having feedback, or information on who had viewed one's location, had a significant impact on how comfortable people were with sharing their information [144]. Burghardt *et al.* went further by exposing individuals to five different privacy technologies in a real world deployment. They reported findings related to both subjects' preferences among the different technologies, and the effectiveness of the technologies [40]. The findings of these two studies are similar to ours, in that they suggest users have rich location-privacy preferences; however, they did not capture these preferences in as much detail as we have done. For example, Burghardt *et al.* asked subjects, prior to being tracked, to report locations that they did not want to share with different groups of individuals (they were given the opportunity to change their reported preferences throughout the study, but were not required to do so). The paper reports some analysis of these privacy preferences suggesting that they are complex. However, the preference collection method used is less detailed than ours and is also somewhat problematic given Connelly *et al.*'s findings [52] that subjects tended to have significant differences between previously asserted and in situ privacy preferences.

The fact that more complex privacy and security settings are needed to capture people's preferences has been observed in other domains as well. For example, Mazurek *et al.* observed that people needed fine-grained access control for configuring their file-sharing preferences [95].

Chapter 7

Conclusion

Mechanisms that facilitate the interactions people have with businesses, their governments, and each other are present everywhere in today's society. One emerging trend over the past decade is a demand for higher levels of expressiveness in such mechanisms. A driving force behind this trend is that greater expressiveness begets better matches, or greater efficiency of the outcomes. Yet, expressiveness does not come for free; it burdens users to specify more preference information. Today's mechanism designers have largely relied on empirical tweaking to determine how to deal with this and related tradeoffs. In this dissertation, we have established the foundation of expressiveness in mechanisms and its relationship to their efficiency, as well as a methodology for determining the most effective forms of expressiveness for a particular setting.

In one stream of research, we proposed a general framework for studying expressiveness in mechanisms based on a novel computational characterization of expressiveness. We showed that the efficiency of an optimally designed mechanism in equilibrium increases *strictly* as we allow more expressiveness. We also showed that, in some cases, a small increase in expressiveness can yield an arbitrarily large increase in a mechanism's efficiency.

In a second stream of research, we operationalized our theory by applying it to a variety of domains. We discussed channel-based mechanisms, which subsume most combinatorial auctions, multi-attribute mechanisms, and the Vickrey-Clarke-Groves scheme. When applied to this class, our general results yield the interesting implication that any (channel-based) multi-item auction that does not allow rich combinatorial bids can be arbitrarily inefficient—unless agents have no private information.

We further operationalized our theory by examining the cost of inexpressiveness in advertisement markets. Using simulated advertiser preferences, we found that, in some realistic settings, slightly increasing the expressiveness of existing ad auction mechanisms leads to significant improvements in their best-case efficiency. The algorithms we discussed for calculating the upper bound on efficiency of different ad auction mechanisms can easily be adapted to other domains, as can our methodology for determining the most appropriate forms of expressiveness (i.e., by simulating the performance different mechanisms under a wide range of preference distributions).

Next, we applied our methodology to the domain of privacy. We discussed an extensive user study that we performed in the context of a location-sharing application. Our study

allowed us to answer questions regarding the most appropriate forms of expressiveness in this context, show how mechanisms can be designed to match the expressions needed by a particular user population, and study the tradeoff between increased accuracy (or efficiency) and the added user burden associated with increased expressiveness.

We concluded with a third application area of catalog pricing. We introduced a framework for automatically mining purchase data and suggesting profit-maximizing prices, bundles, and discounts in more expressive catalog mechanisms. The framework uses a pricing algorithm to compute high-profit prices on items and some bundles, and a model fitting algorithm to estimate a customer valuation model. Our experiments with this framework yielded an interesting finding: in contrast to the suggestions of recommender systems, the most profitable products to offer bundle discounts on appear to be those that are only occasionally purchased together and often separately. We conservatively estimated that a seller with shopping cart data like that of a classic generator, who already has optimally priced items, could increase profits by almost 3% and surplus by over 8% using only bundles of size two (even if he has a thousand items for sale).

All together, the work in this dissertation lends strong support to the thesis that it is possible to improve the efficiency of a wide variety of social and economic mechanisms, in theory and in practice, by using a computational framework for designing them with the most appropriate levels and forms of expressiveness.

7.1 Review of contributions

In Chapter 2, we began by proposing a new **theoretical framework** [20, 23] that characterizes the impact of a mechanism’s expressiveness on its outcome in a domain-independent manner. As part of this work, we introduced two new notions of expressiveness, impact dimension and outcome shattering, based on ideas from computational learning theory. Our main results, such as Theorems 1, 2, and 3, state that a mechanism designer can strictly increase expected efficiency by allowing any agent more expressiveness (until reaching full efficiency). Furthermore, we proved that this can be accomplished with a budget-balanced, Bayes-Nash incentive compatible mechanism (where participants are incentivized to reveal their true valuations in expectation), but we also showed that, without full expressiveness,

it cannot always be accomplished with a mechanism that is dominant-strategy incentive compatible (where participants are incentivized to reveal their true preferences no matter what). We then applied this general framework to a specific class of mechanisms, which we call channel based, and showed that any (channel-based) multi-item auction without rich combinatorial bids can be arbitrarily inefficient.

In the remainder of the dissertation, we operationalized our theoretical framework by developing a **methodology** to compare mechanisms with different degrees and forms of expressiveness in different application domains. At a high level, the methodology, which uses a variety of models, algorithms, and techniques, involves i) estimating preference distributions for participants in a target domain, ii) identifying mechanisms that represent different degrees and forms of expressiveness, iii) computing socially optimal, equilibrium, or heuristic strategies for the agents under each of the mechanisms, iv) simulating the outcomes under the strategies that were computed, and v) comparing the outcomes based on, for example, their expected efficiency.

The first application area we explored (Chapter 3) was that of **advertisement markets** [21]. These markets account for over \$200 billion in annual revenue across all media, and involve some of the fastest-growing mechanisms on the Internet. The most popular online ad mechanism, the generalized second price (GSP) mechanism used by Google, Yahoo!, Bing, Baidu, and others, solicits a single bid from each advertiser for a particular keyword, and assigns advertisers to positions on search-result pages according to these bids. We proved that, since it does not allow advertisers to express different bids for different positions, the GSP is inexpressive according to our domain-independent notions of expressiveness and, consequently, can be arbitrarily inefficient for some preference distributions. However, we also proposed a new mechanism, called the Premium GSP (PGSP), which involves a small, intuitive increase in expressiveness by soliciting a single extra bid from each advertiser (the extra bid is for the right to appear in a premium position). Our empirical results, which involve simulating cooperative and heuristic strategies for the bidders, demonstrated that the PGSP can remove the bulk of the GSP's inefficiency in many realistic settings, which can be up to 30%. Concurrent with our work, Google adopted a feature similar to our premium mechanism, called position preference, suggesting that this type of mechanism is also useful in practice.

The second application area we considered (Chapter 4) was **privacy** [19, 85, 116]. The

past few years have seen an explosion in the range of websites allowing individuals to exchange personal information and content that they have created. These sites include location-sharing services, social-networking services, and photo- and video-sharing services. While there is clearly a demand for people to share this information with each other, there is also a substantial demand for greater expressiveness in the privacy mechanisms that control how the information is shared. To apply our methodology in this domain, we performed a three-week user study in which we tracked the locations of 27 subjects and asked them to rate when, where, and with whom they would have been comfortable sharing their locations. Using the detailed preferences we collected, we identified the best possible policy (or collection of rules granting access to one's location) for each subject and privacy mechanism. To quantify the effects of different levels and forms of expressiveness, we measured the accuracy with which the resulting policies were able to capture our subjects' preferences. We also varied our assumptions about the sensitivity of the information and users' tolerance for the added burden associated with making more complex policies. Our results reveal that many of today's location-sharing applications, such as Loopt and Google's Latitude, may have failed to gain traction due to their limited privacy settings.

In Chapter 5, we investigated a third and final application area of **catalog pricing** [22]. Business to customer retail sales account for nearly four trillion dollars in the United States annually, and the percentage of this shopping done online increased more than three-fold from 2002 to 2007. Yet, despite the increased computational power, connectivity, and data available today, most online and brick-and-mortar retail mechanisms remain nearly identical to their centuries-old original form (i.e., catalog pricing with take-it-or-leave-it offers). This is the default mechanism for brick-and-mortar B2C trade and is used by massive online retailers like Amazon, BestBuy, and Dell. In the final chapter of this dissertation, we began to develop advances toward more expressive catalog pricing mechanisms that could thus lead to significant efficiency improvements across the economy. First, we showed that our theoretical framework for studying expressiveness can be used to characterize the inefficiency of a commonly used inexpressive mechanism: the item-only catalog (i.e., a traditional catalog that offers prices for individual items only). We then described a set of general algorithms for identifying profit-maximizing prices that repeatedly query a customer demand distribution with different candidate catalogs. We provided a method for learning this demand distribution from data, a task that we showed is similar to the classic market basket analysis

problem. (Market basket analysis involves counting the frequencies of different item sets and has been extensively studied). Finally, we performed computational experiments using our pricing and fitting algorithms to demonstrate several conditions under which offering discounts on bundles can benefit the seller, the buyer, and the economy as a whole.

7.2 Review of future work

Our work in Chapter 2 left open several opportunities for future research. First, there is much left to study within channel-based mechanisms, as we discussed at the end of the chapter. We also believe that the efficiency bound and expressiveness measures can be used to provide a richer view of the flaws of inexpressive mechanisms in a wide variety of domains (of course this only provides a bound on this loss for a given mechanism, to compute the loss exactly we would need to extend our analysis to consider the mechanism’s actual equilibrium). In another direction, we can develop algorithms that take as input the prior over the agents’ types in the particular setting at hand and output the efficiency-maximizing mechanism subject to a limit on expressiveness. This objective can be pushed even further to develop a methodology for identifying ways in which existing inexpressive mechanisms can be made more expressive to garner the greatest efficiency increase. Finally, it has often been observed in practice that increases in expressiveness lead to increases in user burden because the increase in expressiveness is typically associated with an increase in the number and/or complexity of “queries” the user has to answer. However, more expressive mechanisms typically eliminate much (or all) of the strategic complexity (e.g., the cognitive effort required to speculate and counter-speculate about the strategies of other agents) that arises when agents are forced to shoehorn their preferences into an inexpressive mechanism. It may be possible to extend our theoretical framework to capture this tradeoff and explore the relationship between these two types of complexity in a variety of settings.

One obvious future research opportunity stemming from our work in Chapter 3 involves using real bidder preference data for ads placed in different positions. However, due to the difficulty in obtaining data about preferences and conversions, it will likely be necessary to adapt our methodology to incorporate other meaningful ways of measuring the inefficiency. For example, rather than relying on real preference data to entirely replace the simulated distributions, one will likely need to develop a “hybrid” distribution that is still partially

simulated, but is more directly informed by real-world data than those we described. It would also be interesting to consider how other types of expressiveness could benefit the GSP. For example, expressions that allow advertisers to bid higher for certain types of users that are likely to convert (e.g., “premium” users). Another future direction is to adapt recent methods for computing equilibria in sponsored search mechanisms by modeling them as action graph games [91, 141] to compute equilibria for our PGSP mechanism. The methodology we have developed can also be adapted and extended to other application domains, such as combinatorial auctions and voting mechanisms.

The findings in Chapter 4 also open several avenues for future work. One avenue involves exploring additional dimensions of privacy preferences. For example, we can study mechanisms that allow users to control the resolution at which location information is provided (e.g., neighborhood, city, or state), or that grant access based on the user’s proximity to the requester. We can also investigate the impact of accuracy models that are richer in terms of their tolerance for error. For example, we can use models with costs for mistakenly revealing a location that depend on the subject, the requester, the time of day, or the location in question. We examined the impact of a rule limit on the accuracy of more expressive privacy mechanisms, but we still assumed that users would be able to identify the most accurate possible rules subject to this limit. This opens up another avenue for future work: accounting for additional cognitive limitations, such as bounded rationality [136], to address issues that challenge this assumption. One potential method for accomplishing this would be to study the behavior of real users of a location-sharing application that offers all of the different expression types discussed in this chapter, such as Locaccino. In such a study we could provide actual users with different privacy mechanisms and measure the amount of sharing that occurs under each. Another interesting aspect to consider in future work is the value of “negative information.” For example, a user who shares his or her location everywhere other than at home is implicitly sharing it at all times, since a requester can infer from a denied request that the user is at home. Finally, there are also legal and policy implications for our work. For example, the information that is protected by the mechanism from other users may not be protected from certain legal entities. In this case, there are two approaches: either the privacy mechanism can be placed on the tracking device itself, which would prevent the information from being recorded, or policies can be enforced that purge the stored data on a regular basis.

Future research stemming from Chapter 5 includes less conservative methods for selecting pricable bundles, discounting bundles of more than two items, and live experiments where the catalogs that we offer serve as demand queries about the customers' valuations that are then incorporated back into our model. These experiments could be carried out similarly to the ones described by Jedidi *et al.* [81], but would involve actual purchases by subjects rather than survey data. There are also several assumptions made in that chapter that could be relaxed in future work. For example, we assumed that each purchase in the shopping cart data was independent, but it may be possible to develop a model that captures repeat purchases by the same customers. We also assumed that the cost of selling an item could be described by a marginal unit cost. It would be interesting to extend our work here to include considerations of non-linear cost functions (e.g., with large start-up costs) or limited-inventory items. Finally, we assumed that the true customer valuations were drawn from distributions that could be accurately fit by our valuation model. However, it would be interesting to consider the effects of mis-representing these valuations because, for example, they are drawn from a different kind of distribution than the one we use (e.g., log normal instead of normal).

Bibliography

- [1] Zoe Abrams, Arpita Ghosh, and Erik Vee. Cost of conciseness in sponsored search auctions. In *Proceedings of the Workshop on Internet Economics (WINE)*, 2007.
- [2] William James Adams and Janet L. Yellen. Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, 90(3):475–498, Aug 1976.
- [3] Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for multi-product pricing. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 72–83, 2004.
- [4] Rakesh Agrawal, Tomasz Imielinski, and Arun Swami. Mining association rules between sets of items in large databases. In *Proceedings of ACM International Conference on Management of Data (SIGMOD)*, pages 207–216, 1993.
- [5] Rakesh Agrawal and Ranakrishnan Srikant. Fast algorithms for mining association rules. In *Proceedings of International Conference on Very Large Data Bases (VLDB)*, 1994.
- [6] Chris Anderson. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, July 2006.
- [7] Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica*, 64(1):51–75, Jan 1996.
- [8] Kenneth Arrow. The property rights doctrine and demand revelation under incomplete information. In M Boskin, editor, *Economics and Human Welfare*. New York Academic Press, 1979.

- [9] Susan Athey and Glenn Ellison. Position auctions with consumer search. *Quarterly Journal of Economics*, 126(3), 2011. Forthcoming.
- [10] Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, Dec 1999.
- [11] Maria-Florina Balcan and Avrim Blum. Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 29–35, Ann Arbor, Michigan, 2006.
- [12] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Item pricing for revenue maximization. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 50–59, Chicago, IL, 2008.
- [13] L. Barkhuus, B. Brown, M. Bell, M. Hall, S. Sherwood, and M. Chalmers. From awareness to repartee: Sharing location within social groups. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2008.
- [14] L. Barkhuus and A. Dey. Location-based services for mobile telephony: A study of users’ privacy concerns. In *Proceedings of the International Conference on Human-Computer Interaction (INTERACT)*, 2003.
- [15] Michael Baye and John Morgan. Brand and price advertising in online markets. Technical Report 0504009, EconWPA, 2005.
- [16] Eyal Beigman and Rakesh Vohra. Learning from revealed preference. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, 2006.
- [17] Shai Ben-David, Nicolo Cesa-Bianchi, David Haussler, and Philip M. Long. Characterizations of learnability for classes of $\{0, \dots, n\}$ -valued functions. *Journal of Computer and Systems Sciences*, 50(1):74–86, 1995.
- [18] Michael Benisch, James Andrews, and Norman Sadeh. Adaptive pricing for customers with probabilistic valuations. In *Revised selected papers from the Proceedings of the AMEC-06 Workshop*, volume 4452 of *Springer Lecture Notes in Computer Science (LNCS)*, pages 132–148, 2007. Earlier version: Proceedings of the International Conference on Electronic Commerce (ICEC) 2006, pp. 38-46.

- [19] Michael Benisch, Patrick Kelley, Norman Sadeh, and Lorrie Cranor. Capturing location-privacy preferences: Quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, pages 1–16, 2010.
- [20] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. A theory of expressiveness in mechanisms. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2008. Extended version: CMU tech report CMU-CS-07-178, 2007.
- [21] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. Methodology for designing reasonably expressive mechanisms with application to ad auctions. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI)*, 2009. Earlier version in proceedings of ACM EC-08 Workshop on Advertisement Auctions.
- [22] Michael Benisch and Tuomas Sandholm. A framework for automated bundling and pricing using purchase data. Mimeo, 2010.
- [23] Michael Benisch and Tuomas Sandholm. A theory of expressiveness in mechanisms. Mimeo, 2010.
- [24] B. Douglas Bernheim and Michael D. Whinston. Menu auctions, resource allocation, and economic influence. *Quarterly Journal of Economics*, 101(1):1–32, 1986.
- [25] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu’alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006.
- [26] Avrim Blum, Jeffrey Jackson, Tuomas Sandholm, and Martin Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004. Early version in COLT-03.
- [27] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):865–929, 1989.
- [28] Liad Blumrosen and Michal Feldman. Implementation with a bounded action space. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 62–71, Ann Arbor, Michigan, USA, 2006.

- [29] Liad Blumrosen, Jason Hartline, and Shuzhen Nong. Position auctions and non-uniform conversion rates. In *Proceedings of the Advertisement Auctions Workshop*, 2008.
- [30] Craig Boutilier. A POMDP formulation of preference elicitation problems. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 239–246, Edmonton, Canada, 2002.
- [31] Craig Boutilier, Ronen Brafman, Christopher Geib, and David Poole. A constraint-based approach to preference elicitation and decision making. In *AAAI Spring Symposium on Qualitative Decision Theory*, 1997.
- [32] Craig Boutilier and Holger H. Hoos. Bidding languages for combinatorial auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.
- [33] Craig Boutilier, David Parkes, Tuomas Sandholm, and William Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2008.
- [34] Craig Boutilier, Tuomas Sandholm, and Robert Shields. Eliciting bid taker non-price preferences in (combinatorial) auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 204–211, San Jose, CA, 2004.
- [35] Sergey Brin, Rajeev Motwani, Jeffrey D. Ullman, and Shalom Tsur. Dynamic itemset counting and implication rules for market basket data. In *Proceedings of the ACM International Conference on Management of Data (SIGMOD)*, pages 255–264, 1997.
- [36] Sergey Brin and Lawrence Page. Dynamic data mining: Exploring large rule spaces by sampling. Mimeo, 1998.
- [37] Christopher H. Brooks and Edmund H. Durfee. Toward automated pricing and bundling of information goods. In *AAAI Workshop on Knowledge-based Electronic Markets*, pages 8–13, 2000.
- [38] N. Brooks. The Atlas rank report: How search engine rank impacts traffic. Available at <http://atlassolutions.com/> in May 2008.

- [39] Zmbl Bulut, lk Grler, and Alper Sen. Bundle pricing of inventories with stochastic demand. *European Journal of Operational Research*, 197(3):897–911, 2009.
- [40] Thorben Burghardt, Erik Buchmann, Jens Müller, and Klemens Böhm. Understanding user preferences and awareness: Privacy mechanisms in location-based services. In *Proceedings of the OnTheMove Conferences (OTM)*, 2009.
- [41] Estelle Cantillon and Martin Pesendorfer. Combination bidding in multi-unit auctions. CEPR Discussion Papers 6083, C.E.P.R. Discussion Papers, 2007.
- [42] K. Chernomaz and D. Levin. Efficiency and synergy in a multi-unit auction with and without package bidding: An experimental study. Mimeo, 2009.
- [43] Chenghuan Sean Chu, Phillip Leslie, and Alan Sorensen. Nearly optimal pricing for multiproduct firms. NBER Working Paper 13916, 2008.
- [44] Ed H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.
- [45] Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, Edmonton, Canada, 2002.
- [46] Vincent Conitzer and Tuomas Sandholm. Applications of automated mechanism design. In *UAI-03 workshop on Bayesian Modeling Applications*, Acapulco, Mexico, 2003.
- [47] Vincent Conitzer and Tuomas Sandholm. Automated mechanism design: Complexity results stemming from the single-agent setting. In *Proceedings of the 5th International Conference on Electronic Commerce (ICEC-03)*, pages 17–24, Pittsburgh, PA, USA, 2003.
- [48] Vincent Conitzer and Tuomas Sandholm. An algorithm for automatically designing deterministic mechanisms without payments. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 128–135, New York, NY, 2004.
- [49] Vincent Conitzer and Tuomas Sandholm. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 132–141, New York, NY, 2004.

- [50] Vincent Conitzer and Tuomas Sandholm. Incremental mechanism design. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, Hyderabad, India, 2007.
- [51] Vincent Conitzer, Tuomas Sandholm, and Paolo Santi. Combinatorial auctions with k -wise dependent valuations. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 248–254, Pittsburgh, PA, 2005. Draft in Oct., 2003.
- [52] Kay Connelly, Ashraf Khalil, and Yong Liu. Do I do what I say?: Observed versus stated privacy preferences. In *Proceedings of the International Conference on Human-Computer Interaction (INTERACT)*, 2007.
- [53] S. Consolovo, I. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: Why, when, & what people want to share. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2005.
- [54] Jason Cornwell, Ian Fette, Gary Hsieh, Madhu Prabaker, Jinghai Rao, Karen Tang, Kami Vaniea, Lujio Bauer, Lorrie Cranor, Jason Hong, Bruce McLaren, Mike Reiter, and Norman Sadeh. User-controllable security and privacy for pervasive computing. In *Proceedings of the 8th IEEE Workshop on Mobile Computing Systems and Applications*, 2007.
- [55] Peter Cramton, Yoav Shoham, and Richard Steinberg. *Combinatorial Auctions*. MIT Press, 2006.
- [56] Nick Craswell, Onno Zoeter, and Michael Taylor. An experimental comparison of click position-bias models. In *Proceedings of the International Conference on Web Search and Data Mining (WSDM)*, 2008.
- [57] Robert E. Dansby and Cecilia Conrad. Commodity bundling. *The American Economic Review*, 74(2):377–381, May 1984.
- [58] Claude d’Aspremont and Louis-Andre Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

- [59] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *The American Economic Review*, 97(1):242–259, March 2007.
- [60] Eyal Even-Dar, Michael Kearns, and Jennifer Wortman. Sponsored search with contexts. In *Proceedings of the Workshop on Internet Economics (WINE)*, 2007.
- [61] G. Gallego and G. van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations Research*, 45(1):24–41, 1997.
- [62] Sorabh Gandhi, Chiranjeeb Buragohain, Lili Cao, Haitao Zheng, and Subhash Suri. A general framework for clearing auctions of wireless spectrum. In *Proceedings of IEEE Dynamic Spectrum Access Networks (DySPAN)*, 2007.
- [63] Marta C. Gonzalez, Cesar A. Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.
- [64] Kelsey Group. BIA’s The Kelsey Group Forecasts U.S. Mobile Local Search Advertising Revenues to reach \$1.3B in 2013, February 2009. <http://www.kelseygroup.com/press/>.
- [65] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [66] Joseph P. Gultinan. The price bundling of services: A normative framework. *Journal of Marketing*, 51(2):74–85, Apr 1987.
- [67] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On profit-maximizing envy-free pricing. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.
- [68] MohammadTaghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2007.
- [69] Ward Hanson and R. Kipp Martin. Optimal bundle pricing. *Management Science*, 36(2):155–174, 1990.

- [70] Friedrich Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, 1945.
- [71] Jeffrey Hightower, Anthony LaMarca, and Ian E. Smith. Practical lessons from Place Lab. *IEEE Pervasive Computing*, 5(3):32–39, 2006.
- [72] Lorin M. Hitt and Pei-yu Chen. Bundling with customer self-selection: A simple approach to bundling low-marginal-cost goods. *Management Science*, 51(10):1481–1493, 2005.
- [73] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars, Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.
- [74] Ron Holzman, Noa Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, April 2004.
- [75] S. Huang, F. Proulx, and C. Ratti. iFIND: a Peer-to-Peer application for real-time location monitoring on the MIT campus. In *International Conference on Computers in Urban Planning and Urban Management (CUPUM)*, 2007.
- [76] Leonid Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*. Amsterdam: North Holland, 1972.
- [77] Leonid Hurwicz. On informationally decentralized systems. In C.B McGuire and R. Radner, editors, *Decision and Organization*, chapter 14, pages 297–336. University of Minnesota Press, 1986. Second edition.
- [78] Leonid Hurwicz and Thomas Marschak. Discrete allocation mechanisms: Dimensional requirements for resource-allocation mechanisms when desired outcomes are unbounded. *Journal of Complexity*, 1(2):264–303, 1985.
- [79] Leonid Hurwicz and Stanley Reiter. *Designing Economic Mechanisms*. Number 9780521724104 in Cambridge Books. Cambridge University Press, 2006.

- [80] G. Iachello, I. Smith, S. Consolovo, G. Abowd, J. Hughes, J. Howard, F. Potter, J. Scott, T. Sohn, J. Hightower, and A. LaMarca. Control, deception, and communication: Evaluating the deployment of a location-enhanced messaging service. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2005.
- [81] Kamel Jedidi, Sharan Jagpal, and Puneet Manchanda. Measuring heterogeneous reservation prices for product bundles. *Marketing Science*, 22(1):107–130, 2003.
- [82] Philippe Jehiel, Moritz Meyer-Ter-Vehn, and Benny Moldovanu. Mixed bundling auctions. *Journal of Economic Theory*, 127(1):494–512, 2007.
- [83] Philippe Jehiel and Benny Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, Sept 2001.
- [84] Przemyslaw Jeziorski and Ilya Segal. What makes them click: Empirical analysis of consumer demand for internet search advertising. Mimeo, June 2009.
- [85] Patrick Kelley, Michael Benisch, Lorrie Cranor, and Norman Sadeh. When are users comfortable sharing locations with advertisers? In *Proceedings of the Conference on Human Factors in Computing (CHI)*, 2010. Note. Earlier version: Carnegie Mellon University Technical Report CMU-ISR-10-126.
- [86] Jeffrey O. Kephart, Christopher H. Brooks, Rajarshi Das, Jeffrey K. MacKie-Mason, Robert Gazzale, and Edmund H. Durfee. Pricing information bundles in a dynamic environment. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 180–190, Tampa, Florida, 2001.
- [87] E Kushilevitz and N Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [88] Sebastián Lahaie and David Parkes. Applying learning algorithms to preference elicitation. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, New York, NY, 2004.
- [89] Sébastien Lahaie, David C. Parkes, and David M. Pennock. An expressive auction design for online display advertising. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2008.

- [90] S. Lederer, J. Mankoff, and A. K. Dey. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2003.
- [91] Kevin Leyton-Brown and Moshe Tennenholtz. Local-effect games. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, Acapulco, Mexico, 2003.
- [92] Anton Likhodedov and Tuomas Sandholm. Mechanism for optimally trading off revenue and efficiency in multi-unit auctions. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, New York, NY, 2004. Short paper. Earlier versions in AMEC-03 and ACM-EC-03.
- [93] F. Marechal and P. H. Morand. Free riding in combinatorial first-price sealed-bid auctions. *The B. E. Journal of Theoretical Economics*, 9(1), 2009.
- [94] Andreu Mas-Colell, Michael Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [95] M. Mazurek, J.P. Arsenault, J. Bresee, N. Gupta, I. Ion, C. Johns, D. Lee, Y. Liang, J. Olsen, B. Salmon, R. Shay, K. Vaniea, L. Bauer, L.F. Cranor, G.R. Ganger, and M.K. Reiter. Access control for home data sharing: Attitudes, needs and practices. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2010.
- [96] R. Preston McAfee, John McMillan, and Michael D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, 104(2):371–383, May 1989.
- [97] Theresa Metty, Rob Harlan, Quentin Samelson, Tom Moore, Thomas Morris, Ron Sorensen, Avner Schneur, Olga Raskina, Rina Schneur, Joshua Kanner, Kevin Potts, and Jeffrey Robbins. Reinventing the supplier negotiation process at Motorola. *Interfaces*, 35(1):7–23, 2005.
- [98] Paul Milgrom. Simplified mechanisms with applications to sponsored search and package auctions. *Games and Economic Behavior*, 70:62–70, 2010.

- [99] Claire Cain Miller and Jenna Wortham. Technology aside, most people still decline to be located. <http://www.nytimes.com/2010/08/30/technology/30location.html>.
- [100] Tom Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [101] Kenneth Mount and Stanley Reiter. The informational size of message spaces. *Journal of Economic Theory*, 8(2):161–192, 1974.
- [102] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.
- [103] Victor Naroditsky and Amy Greenwald. Using iterated best-response to find Bayes-Nash equilibria in auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Vancouver, Canada, 2007. Short paper.
- [104] Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.
- [105] Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129:192–224, 2006.
- [106] M. Olivares, G. Y. Weintraub, D. Yung, and R. Epstein. The design of combinatorial auctions for procurement: An empirical study of the chilean school meals auction. Mimeo, 2009.
- [107] Shmuel Oren and Michael Rothkopf. Optimal bidding in sequential auctions. *Operations Research*, 23:1080–1090, 1975.
- [108] Michael Ostrovsky and Michael Schwarz. Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the Advertisement Auctions Workshop*, 2010.
- [109] Thomas Palfrey. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica*, 51:463–484, 1983.
- [110] David Parkes. iBundle: An efficient ascending price bundle auction. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 148–157, Denver, CO, November 1999.

- [111] David Parkes. Auction design with costly preference elicitation. *Annals of Mathematics and Artificial Intelligence*, 44:269–302, 2005. Special Issue on the Foundations of Electronic Commerce.
- [112] David C. Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2001.
- [113] David C. Parkes. Price-based information certificates for minimal-revelation combinatorial auctions. In *AAMAS '02: Revised Papers from the Workshop on Agent Mediated Electronic Commerce on Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems*, pages 103–122, London, UK, 2002. Springer-Verlag.
- [114] David C. Parkes and Tuomas Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *Proceedings of the Workshop on Sponsored Search Auctions*, Vancouver, BC, Canada, June 2005.
- [115] S. Patil and J. Lai. Who gets to know what when: Configuring privacy permissions in an awareness application. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2005.
- [116] Ramprasad Ravichandran, Michael Benisch, Patrick Gage Kelley, and Norman Sadeh. Capturing social networking privacy preferences: Can default policies help alleviate tradeoffs between expressiveness and user burden? In *Proceedings of Privacy Enhancing Technologies (PETS)*, volume 5672 of *Springer Lecture Notes in Computer Science (LNCS)*, pages 1–18, 2009. Earlier version: Symposium on Usable Privacy and Security (SOUPS) 2008, poster.
- [117] Valentin Robu, Han Poutré, and Sander Bohte. *The Complex Dynamics of Sponsored Search Markets*, pages 183–198. Springer-Verlag, Berlin, Heidelberg, 2009.
- [118] Amir Ronen. Mechanism design with incomplete languages. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*, pages 105–114, 2001.
- [119] Robert W. Rosenthal and Ruqu Wang. Simultaneous auctions with synergies and common values. *Games and Economic Behavior*, 17:32–55, 1996.

- [120] Paat Rusmevichientong, Benjamin Van Roy, and Peter W. Glynn. A nonparametric approach to multiproduct pricing. *Operations Research*, 54(1):82–98, 2006.
- [121] Norman Sadeh, Fabien Gandon, and Oh Buyng Kwon. Ambient intelligence: The MyCampus experience. In T. Vasilakos and W. Pedrycz, editors, *Ambient Intelligence and Pervasive Computing*. ArTech House, 2006.
- [122] Norman Sadeh, Jason Hong, Lorrie Cranor, Ian Fette, Patrick Kelley, Madhu Prabaker, and Jinghai Rao. Understanding and capturing people’s privacy policies in a mobile social networking application. *The Journal of Personal and Ubiquitous Computing*, 13(6):401–412, 2009.
- [123] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, January 2002.
- [124] Tuomas Sandholm. Automated mechanism design: A new application area for search algorithms. In *International Conference on Principles and Practice of Constraint Programming (CP)*, pages 19–36, Cork, Ireland, 2003.
- [125] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
- [126] Tuomas Sandholm and Craig Boutilier. Preference elicitation in combinatorial auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, pages 233–263. MIT Press, 2006. Chapter 10.
- [127] Tuomas Sandholm, Vincent Conitzer, and Craig Boutilier. Automated design of multistage mechanisms. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI)*, Hyderabad, India, 2007.
- [128] Tuomas Sandholm and Vincent Ferrandon. Safe exchange planner. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, pages 255–262, Boston, MA, July 2000.

- [129] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2006.
- [130] Tuomas Sandholm, David Levine, Michael Concordia, Paul Martyn, Rick Hughes, Jim Jacobs, and Dennis Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.
- [131] Tuomas Sandholm, Anton Likhodedov, and Andrew Gilpin. Automated design of revenue-maximizing combinatorial auctions. *Operations Research*, 2011. Special issue on Computational Economics. To appear.
- [132] Tuomas Sandholm and Subhash Suri. Side constraints and non-price attributes in markets. *Games and Economic Behavior*, 55:321–330, 2006. Extended version in IJCAI-2001 Workshop on Distributed Constraint Reasoning.
- [133] Richard Schmalensee. Gaussian demand and commodity bundling. *The Journal of Business*, 57(1):S211–S230, Jan 1984.
- [134] Tobias Schoenherr and Vincent A. Marbert. The use of bundling in B2B online reverse auctions. *Journal of Operations Management*, 2008.
- [135] Barry Schwartz. *The Paradox of Choice: Why More is Less*. Ecco/HarperCollins Publishers, 2004.
- [136] Herbert A Simon. *Models of bounded rationality*, volume 2. MIT Press, 1982.
- [137] I. Smith, S. Consolovo, A. LaMarca, J. Hightower, J. Scott, T. Sohn, J. Hughes, G. Iachello, and G. Abowd. Social disclosure of place: From location technology to communication practices. In *Lecture Notes in Computer Science : Pervasive Computing*, pages 134–151, 2005.
- [138] George J. Stigler. United States v. Loew’s Inc.: A note on block-booking. *The Supreme Court Review*, 1963:152–157, 1963.
- [139] Balazs Szentes and Robert W. Rosenthal. Beyond chopsticks: Symmetric equilibria in majority auction games. *Games and Economic Behavior*, 45(2):278–295, 2003.

- [140] Kalyan T. Talluri and Garrett John Van Ryzin. *The theory and practice of revenue management*. Number 68 in International series in operations research and management science. Kluwer Academic Publisher, Boston, Mass. [u.a.], 2004.
- [141] David R.M. Thompson and Kevin Leyton-Brown. Computational analysis of perfect-information position auctions. In *ACM Conference on Electronic Commerce (ACM-EC)*, 2009.
- [142] Eran Toch, Justin Cranshaw, Paul Hanks Drielsma, Janice Y. Tsai, Patrick Gage Kelley, James Springfield, Lorrie Cranor, Jason Hong, and Norman Sadeh. Empirical models of privacy in location sharing. In *Proceedings of the International Conference on Ubiquitous Computing (UbiComp)*, 2010.
- [143] J. Tsai, P. Kelley, L. Cranor, and N. Sadeh. Location-sharing technologies: Privacy risks and controls. *A Journal of Law and Policy for the Information Society (I/S)*, 6(2), 2010.
- [144] Janice Tsai, Patrick Kelley, Paul Hanks Drielsma, Lorrie Faith Cranor, Jason Hong, and Norman Sadeh. Who’s viewed you? the impact of feedback in a mobile-location system. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 2009.
- [145] US Census Bureau. Statistical abstract of the United States: Domestic trade, 2005. http://www.census.gov/prod/www/abs/statab2001_2005.html.
- [146] US Census Bureau. Statistical abstract of the United States, wholesale & retail trade: Online retail sales, 2010. http://www.census.gov/compendia/statab/cats/wholesale_retail_trade/online_retail_sales.html.
- [147] Vladimir Vapnik and Alexey Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.
- [148] H. R. Varian. The non-parametric approach to demand analysis. *Econometrica*, 50:945 – 974, 1982.

- [149] Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, pages 1163–1178, 2007.
- [150] R. Venkatesh and W. Kamakura. Optimal bundling and pricing under a monopoly: Contrasting complements and substitutes from independently valued products. *Journal of Business*, 76(2):211–231, 2003.
- [151] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
- [152] W. Walsh, C. Boutilier, T. Sandholm, R. Shields, G. Nemhauser, and D. Parkes. Automated channel abstraction for advertising auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2010.
- [153] Yi Wang, Jialiu Lin, Murali Annavaram, Quinn A. Jacobson, Jason Hong, Bhaskar Krishnamachari, and Norman Sadeh. A framework of energy efficient mobile sensing for automatic user state recognition. In *Proceedings of the International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2009.
- [154] Roy Want, Veronica Falcão, and Jonathan Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10:91–102, 1992.
- [155] Jim Wilenius and Arne Andersson. Discovering equilibrium strategies for a combinatorial first price auction. In *Proceedings of the E-Commerce Technology Conference (CEC)*, 2007.
- [156] Shin-yi Wu, Lorin M. Hitt, Pei-yu Chen, and G. Anandalingam. Customized bundle pricing for information goods: A nonlinear mixed-integer programming approach. *Management Science*, 54(3):608–622, 2008.
- [157] Manjit S. Yadav and Kent B. Monroe. How buyers perceive savings in a bundle price: An examination of a bundle’s transaction value. *Journal of Marketing Research*, 30(3):350–358, Aug 1993.
- [158] A. C. Yao. Some complexity questions related to distributed computing. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 209–213, 1979.

- [159] Weili Zhu and Peter Wurman. Structural leverage and fictitious play in sequential auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Edmonton, Canada, 2002.

