

Overcoming Incentive Constraints

by Matthew O. Jackson and Hugo F. Sonnenschein *

This Draft: November 24, 2003[†]

Abstract

Consider an arbitrary Bayesian decision problem in which the preferences of each agent are private information. We prove that the utility costs associated with incentive constraints typically decrease when the decision problem is linked with independent copies of itself. This is established by first defining a mechanism in which agents must budget their representations of preferences so that the frequency of preferences across problems mirrors the underlying distribution of preferences, and then arguing that agents will satisfy their budget by being as truthful as possible. Examples illustrate the disappearance of incentive costs when problems are linked in a rich variety of problems, including public goods allocation, voting, and bargaining.

*Jackson is at the Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, California 91125, USA, <http://www.hss.caltech.edu/~jacksonm/Jackson.html>, email: jacksonm@hss.caltech.edu, and Sonnenschein is at the Department of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637, USA, email: h-sonnenschein@uchicago.edu. Financial support under NSF grants SES-9986190 and SES-0316493 is gratefully acknowledged, as is support from the University of Chicago where Jackson was visiting when this research began. We thank Zachary Cohn and Phil Reny for calling our attention to the treatment of participation constraints, Bill Zame for questions about coalitional considerations, Mahmut Erdem for written suggestions on an earlier draft, Robert Veszteg as a discussant of the paper, and the participants of various seminars for helpful questions and comments.

[†]This supercedes an earlier version of this paper entitled “The Linking of Collective Decisions and Efficiency,” Caltech Social Science Working Paper 1159, March 2003.

1 Introduction

Over the past fifty years we have learned that social welfare possibilities depend not only on resources and technology, but equally and most critically on incentive constraints, including participation constraints, and the ability of social institutions to mediate those constraints. Thus, voting systems, labor contracts, financial contracts, auction forms, and a host of other practical arrangements are now commonly formulated as Bayesian games, and judged in terms of their ability to mediate incentive constraints.

This paper demonstrates that the limitations that incentive constraints impose on the attainment of socially efficient outcomes can be made to disappear when several problems are linked. We exploit the idea that when several independent social decision problems are linked, or when there are several independent aspects of a given problem, then it makes sense to speak of “rationing” or “budgeting” an agent’s representations. As in everyday experience, agents may be asked to “reveal the issues that they regard as most important”, and the position that one’s “needs are extreme with respect to all aspects of the offer that is on the table” may be taken as a signal that an agent is not serious about a negotiation. Here, when agents are asked to reveal their preferences over different problems or different aspects of a problem, they are not permitted to claim to have extreme preferences over all of them. For instance, when a buyer and seller bargain over several items, we do not allow the buyer to claim to have a low valuation for each item, nor do we allow the seller to claim to have a high cost for all of the items. The critical insight, is that by linking the bargaining over the items, we can discover which items are relatively more costly for the seller and which items are relatively more valued by the buyer. The rationing or budgeting of announcements leads to a tendency, which we make precise, for agents to be as truthful in their representation as is possible. This helps us to overcome incentive constraints and yields more efficient decisions.

In more formal language, we consider an abstract Bayesian collective decision problem and an ex ante Pareto efficient social choice function f that indicates the collective decision we would like to make as a function of the realized preferences of the n agents. Let (u_1, u_2, \dots, u_n) denote the ex ante expected utilities that are achieved under f . Such an ideal f will generally not be implementable because of incentive constraints. Now, consider K copies of the decision problem, where agents’ preferences are additively separable and independently distributed across the problems. We show that as K becomes large it is possible to essentially implement f on each problem and thus achieve the target utilities (u_1, u_2, \dots, u_n) on each problem. Even when K is small there are typically substantial utility gains from considering the K problems together.

We establish this result by constructing a general mechanism that has each agent present a K -vector of preferences, and then the decision on each of the K problems is made according to f . The key is that we require that agents present vectors whose distribution of types across problems mirrors the underlying distribution of their preferences. The agents are not, for example, allowed to represent themselves as having a "bad draw" on more than the expected number of problems on which they "should" have a bad draw. Each agent is forced to budget her representations, including her lies, in such a manner that they do her the most good. With some adjustments in this idea to deal with the case of more than two agents, as well as participation constraints, we show that in the limit there is no gain from lying. In fact, for every K there is an equilibrium in which all agents are telling the truth as fully as the constraint on their representations permits. Moreover, when all other agents mirror the underlying distribution of preferences (as forced to by the mechanism), being as truthful as possible secures for the remaining agent an expected utility that for large K approximates the target utility level regardless of the strategies followed by the other agents. Thus, we can conclude that *all* equilibria of the linking mechanisms converge to the target utility levels.

We should emphasize that while we make use of the law of large numbers, the heart of the matter is to figure out who is of what type on which problems. A large number of independent problems makes it likely that the realized frequency of agents' preferences mirrors the underlying distribution of preferences, but this cannot guarantee that agents will find it in their interest to truthfully announce those preferences. It is the interplay between the efficiency of the underlying social choice function and the fact that agents' announcements are budgeted to match the underlying distribution that makes it in each agent's best interest (in a strong sense that we make precise) to be as truthful as she can. Indeed, in a technical sense, this is what the paper is all about.

We believe that it is useful to explain our ideas by first considering a variety of examples in which incentive constraints prevent the achievement of a first best optimum, and then, in the context of these examples, demonstrating how and why linking several problems can increase per problem efficiency. The problems have different structures and the lessons that are accumulated suggest the general treatment. These examples also demonstrate that significant efficiency gains are possible from linking even a few problems, and show the failure of existing mechanisms (e.g., Clarke-Groves-Vickrey mechanisms, or variations on those due to d'Aspremont and Gerard-Varet (1979)) to achieve the same implementation. Furthermore, we believe that an exposition that begins with the general abstract problem might make it difficult for the reader to appreciate the intimate connection that our work has to classical issues in voting theory, the provision of public goods, bargaining theory, and

indeed virtually all settings in which incentive or agency costs play a prominent role.

Finally, we come back to discuss the related literature at the end of our paper. The fact that one can find mechanisms scattered in very different literatures that have observed various aspects or possibilities of utility gains from linkage, but not understood this as a general phenomenon, reinforces our approach, as we discuss in Section 6.

2 Examples

Our first example concerns voting over several issues when the preferences of agents are private information.

EXAMPLE 1 *A Voting Problem*

Consider a society that must make a single decision that affects all of its members' well-being. For example, the decision may be whether or not to undertake a given project or law, possibly including a specification of how the costs of the implemented project will be distributed. The society might have an arbitrary number of agents and could be choosing from any number of alternatives. However, for the purpose of this example, consider the case where the decision is binary represented by $d \in \{a, b\}$ and there are just two agents, and we return to the general case later.

The agents have utilities for each possible decision. Let $v_i(d)$ denote agent i 's value for taking decision d . The important information for this binary problem is simply the difference in valuations between decisions a and b . An agent's preferences are thus summarized by the difference in utilities between decisions a and b , denoted $v_i = v_i(b) - v_i(a)$.

If v_i is positive for both agents, then the unanimously preferred decision is $d = b$, and if v_i is negative for both agents, then the unanimously preferred decision is $d = a$. In the case where $v_i > 0$ while $v_j < 0$, then which decision should be made is ambiguous. To keep things simple for now, consider the case where each v_i is independently and identically distributed and takes on values in $\{-2, -1, 1, 2\}$ with equal probability.

Suppose that we wish to choose the decision that maximizes the sum of the utilities, and in the case of a tie we flip a coin. That is, we wish to make the decision as follows:

		Agent 2's value			
		2	1	-1	-2
Agent 1's value	2	b	b	b	coin
	1	b	b	coin	a
	-1	b	coin	a	a
	2	coin	a	a	a

The difficulty is that this social choice function viewed is not incentive compatible. For instance, either agent would be better off saying that his or her type is $v_i = 2$ when it is really $v_i = 1$, regardless of what the other agent does.

The “closest” social choice function which is incentive compatible is the most prominent system for making such binary decisions: simply voting and flipping a coin in the event of a tie.

		Agent 2's value			
		2	1	-1	-2
Agent 1's value	2	b	b	coin	coin
	1	b	b	coin	coin
	-1	coin	coin	a	a
	2	coin	coin	a	a

This social choice function is the unique incentive compatible one that maximizes the total sum of utilities subject and is anonymous and neutral (a version of May's (1952) theorem).¹ It is also ex post Pareto efficient relative to the realized preferences.

However, it is important to note that the simple voting mechanism is not ex ante efficient. This can be seen simply from a comparison of the two tables above. Both agents would prefer to make the following improvements: if the agents disagree and one of the agents has an intensity of 2 while the other has an intensity of 1 but of the opposite signs, then the decision is made in favor of the agent who cares more, rather than being decided by the flip of a coin. This sometimes goes against an agent's wishes and sometimes for the agent. The reason that this improves over flipping a coin is that it goes in the agent's favor in situations where the agent cares more intensely and against the agent when the agent cares less intensely. The big problem with this improvement, of course, is that it is not incentive compatible. If we try to ask the agents whether they care a lot or a little, they are always better off pretending to care a lot.

¹There are other incentive compatible social choice function that are not anonymous and/or neutral, such as constant mechanisms, dictatorships, favoring one alternative over the other in the event of a tie, or variations on such mechanisms.

Linking Two Such Decisions

Next, let us consider a situation where there are two separate decisions to be made. These might even be two different “features” of a single decision. Let us label them $d_1 \in \{a_1, b_1\}$ and $d_2 \in \{a_2, b_2\}$. Here each agent has preferences over each decision, and values a combination (d_1, d_2) according to $v_i(d_1, d_2) = v_{i1}(d_1) + v_{i2}(d_2)$. Again, we can characterize the preferences over a decision d_j by a utility difference $v_{ij} = v_{ij}(b) - v_{ij}(a)$. Let these v_{ij} ’s be i.i.d. on $\{-2, -1, 1, 2\}$ with equal probabilities. Thus, we are considering a duplication of the previous decision problem.

One approach to solving this problem is to hold separate votes over the two problems. Note, however, that this is no longer even ex post efficient! To see this, consider a situation where agent 1 has values (2,1) for the respective problems, and agent 2 has values (-1,-2) for the respective problems. The votes will be tied on both decision problems, and coin flips will decide each. One possible outcome of the coin flips results in a decision of (a_1, b_2) . This outcome is Pareto inefficient as both agents would prefer to have the decision of (b_1, a_2) .

It is useful to note that what was an ex ante inefficiency in the isolated problem, becomes an ex post inefficiency in the duplicated problem. Effectively the trades that agents would like to make across possible states in the isolated problem, become trades that the agents would like to make across different problems in the duplicated setting! This allows us to find mechanisms that do better in the setting with two problems; and in fact, it even offers us a pretty good suggestion as to how we should do this.

Consider the following linked mechanism that operates over the two problems. We ask agents to announce their utilities for each problem. However, we constrain them in the following manner: An agents announced pair of utilities (v_{i1}, v_{i2}) must have one utility of magnitude 2 and one utility of magnitude 1. We then run the target ex ante efficient social choice function, f described earlier, on these constrained announcements. Thus, if agents announcements agree on the sign, we choose the alternative that they both favor. If the agents disagree on sign, then we decide in favor of the agent whose announced utility has a larger magnitude and flip a coin in the event of a tie on magnitudes.²

²Note that we can also implement the outcomes of this linked mechanism in an alternative way. We give agents each three (indivisible) votes and require them to cast at least one vote on each problem. This is reminiscent of Casella’s (2002) storable votes which may be spread across time. However, we have placed more restrictions on votes (requiring that one be spent on each problem) which helps enhance efficiency. Also, once we move beyond this simple two decision-two intensity voting setting our approach bears little relationship to storable votes, as we shall see shortly in examples of public goods settings, bargaining settings, and others.

It is straightforward to check that there is a Bayesian equilibrium of this mechanism with the following features:

- if an agent's magnitude of utility differs across the two problems then he or she announces utilities truthfully
- if an agent has two utilities of the same magnitude, then the agent announces the correct signs but the agent randomly chooses which problem to announce the higher magnitude on.

In fact, all equilibria of this mechanism have similar features up to the tie breaking, something that we will come back to discuss more generally below.

The equilibria of the linked mechanism is not quite *ex ante* Pareto efficient. Nevertheless, the equilibrium outcomes of the linked mechanism still Pareto dominate from any perspective (*ex ante*, *interim*, or *ex post*) voting on the problems separately.³ To get a feeling for the level of Pareto improvement of the linked mechanism over the separate voting, let's look at the probability of not choosing the outcome prescribed by f . The linked mechanism has cut the probability of making such errors in half relative to that of running two separate voting mechanisms. To see this, first note that the only situations where errors can arise are on problems where the agents disagree both on sign and magnitude of preference. Conditional on this case, a separate (non-linked) voting mechanism will flip a coin and make an error with probability $1/2$. In the linked mechanism, an error will only occur with probability $1/4$.⁴ Thus, linking the problems has cut the probability of making an error on each problem in half.

³In *ex post* comparisons one has to be careful about the timing: before or after coin flips. There is a chance that an agent gets lucky and wins all coin flips, and so comparisons after coin flips makes the two mechanism non-comparable. However, if we make comparisons after types are known, but before coins are flipped, then the linked mechanism dominates the separate voting.

⁴This is seen as follows. There are four equally likely sub-cases: (a) each agent's magnitude on the other problem differs from that on this problem, which implies that announcements will be truthful and no error will be made; (b) agent 1 has equal magnitudes across the problems but not agent 2, in which case there is a conditional probability of $1/2$ that the two agents' announcements will match and then a conditional probability of $1/2$ that the coin flip will result in an error - so a probability of $1/4$ of an error conditional on this sub-case; (c) agent 2 has equal magnitudes across the problems but not agent 1; and so this is analogous to sub-case (b); and (d) both agents have equal magnitudes across the problems in which case the announcements and coin flip are all essentially random and the conditional probability of an error is $1/2$. As each sub-case occurs with probability $1/4$, we have a probability of $1/4$ ($\frac{1}{4}(0 + \frac{1}{4} + \frac{1}{4} + \frac{1}{2})$) of making an error in total across the sub-cases.

The reason why linking the two decisions together improves efficiency is as follows. Linking the decisions together allows us to ask a question of the form, “Which decision do you care more about?” This can be answered in an incentive compatible way in the linked problem, but we cannot even ask this question in the separate problem. Effectively, linking the problem has changed things that were ex ante inefficiencies - “I would like to make trades over my different possible future selves,” to ex post inefficiencies - “I now actually have different selves and would be happy to make trades across them”. So fixing ex post inefficiencies in the linked problem, is in a sense overcoming ex ante inefficiencies that could not be overcome in the original problem.

Finally, we observe that no *interpersonal* comparability in utilities is needed in the above analysis. The ex ante inefficiency in straight voting is due to the fact that both agents would be willing to make trades across different states if they could. It is *intrapersonal* comparisons that are at the heart here. All of the points that we make in this paper are valid even if we work with forms of Pareto efficiency that don’t make any implicit *interpersonal* comparisons.

Linking Many Such Decisions

We have seen that linking two decisions together helps improve the total performance of the optimal mechanism. Still, it did not reach complete efficiency. What if we link more decisions together? Indeed, linking more decisions together helps further, and in the limit leads us to full Pareto efficiency.

This is easily seen in the context of the above example. Suppose that we have linked K independent decisions together of the type described above, where K is a “large” number. Consider the following mechanism. The agents cast a vote on each problem j for either a_j or b_j . The agents are also allowed to declare $\frac{K}{2}$ problems for which they care more intensely for; that is, for which $|v_{ij}| = 2$. If there is a tie in the vote, then the tie is broken in favor of the agent who has declared they care more intensely for the problem, if there is exactly one such agent, and otherwise a fair coin is flipped.

With large n , the agents will care intensely for approximately $\frac{1}{2}$ of the problems. They may end up caring intensely for a few more or less problems than exactly $\frac{1}{2}$, in which case the mechanism will force them to “lie” on some small fraction of problems. However, again there exists an equilibrium where agents are always truthful about the signs of their utility for the problems and are truthful about magnitude up to the extent that they can be under the constraints. That is, if an agent cares intensely about more than $\frac{K}{2}$ problems, then the agent randomly picks $\frac{K}{2}$ of those to declare as high magnitude and declares low magnitude on the others; and similarly for the case where an agent has a low magnitude on more than $\frac{K}{2}$ problems.

As K becomes large, the fraction of problems where agents' announcements are not completely truthful goes to 0, and so the probability that the decision on any given problem is incorrect goes to 0. So, on each problem, we are converging to truthful revelation and the associated ex ante (and thus interim and ex post) efficient decisions. Moreover, we shall argue below that *all* equilibria of this mechanism converge to the truthful revelation utilities.⁵

This voting example has provided some of the basic ideas that underlie more general results. We note that in this example we have not considered any transfers or money in the operation of the mechanism. This is an advantage of our approach, in that it works even in settings where one cannot (or prefers not to) use transfers. In the context of the voting example, however, if one could use transfers then that would be another way to mitigate the inefficiency problem in the context of Example 1. Yet, we should emphasize that, as the knowledgeable reader will know, allowing for transfers cannot always reconcile incentive and participation constraints with efficiency. We now present another example, where even in the presence of transfers there is a conflict between efficiency and incentive compatibility and participation constraints. Here we show that linking problems can also achieve approximately efficient outcomes in contexts where transfers are already freely available and are unable to achieve efficient outcomes.⁶

EXAMPLE 2 *A Public Goods Example*

Consider a decision by a society of n agents of whether or not to build a public project. The project costs $c > 0$. Agents have values for the public good that fall in the set $\{0, 1, \dots, m\}$, and are denoted v_i . Let v denote the vector of values. Suppose that this

⁵The linking method we have proposed can be further improved upon along the sequence, by taking advantage of some specific aspects of the problem. Start with voting and declarations of which problems agents care more intensely for, just as above. However, allow an agent to designate more or fewer than $\frac{K}{2}$ problems that they care intensely for, and then let the mechanism choose for the agent on which problems to assign a higher magnitude - so that the number of such announcements still comes out at $\frac{K}{2}$. The mechanism picks these problems by coordinating across the two agents in such a way to best match the announcements. So, each agent still has rights to claim to care intensely about $\frac{K}{2}$ problems. However, when an agent happens to care about fewer problems, in the previous mechanism they would end up picking some extras randomly. It is actually more efficient to coordinate those across agents, so that one agent's "lies" don't fall on problems where the other agent truly cares intensely. By allowing the mechanism instead of the agents to pick the "lies," efficiency is improved. Cohn (2003) examines the magnitudes of potential improvements over our mechanism in the context of Example 3.

⁶It might be interesting to see to what extent judicious use of transfers in conjunction with linking mechanisms might speed the rate of convergence to efficiency that we obtain.

is a transferable utility world, so that Pareto efficiency dictates that we should build the public good when $\sum_i v_i > c$ and not when $\sum_i v_i < c$. Moreover, we suppose that we would like to split the costs among the agents in a way so that no agent's share of the cost exceeds their valuation. So, each agent will pay a cost share $c_i(v)$ such that $c_i(v) \leq v_i$, and $\sum_i c_i(v) = c$ when $\sum_i v_i > c$, and $c_i(v) = 0$ otherwise.

The target decision rule that we have described will in general not be incentive compatible. To see this is straightforward. For instance, take the simple case where $n = 2$, $m = 1$ and $c < 1$. Here, if at least one agent has $v_i = 1$, then we build the project and split the costs equally among those having $v_i = 1$. The target rule can be pictured as follows.

The Target (Ex Ante Efficient) Outcomes

		Agent 2's Value	
		1	0
Agent 1's	1	Build, Agent 1's Cost = $c/2$	Build, Agent 1's Cost = c
Value	0	Build, Agent 1's Cost = 0	Not Build

Let us check that this is not incentive compatible. Consider a case where the probability of having a valuation for the project of $v_i = 1$ is $\frac{2}{3}$ and the probability of having a valuation of $v_i = 0$ is $\frac{1}{3}$. Consider an agent who has a valuation of $v_i = 1$. By pretending to have $v_i = 0$, and supposing that the other agent is truthful about his or her value, the given agent will still enjoy the public project with probability $\frac{2}{3}$, but save on paying the cost. This comes at some risk, as pretending to have $v_i = 0$ may result in not having the project built if it turns out that the other agent has a valuation of 0, which happens with probability $\frac{1}{3}$. In particular the overall expected cost savings is $\frac{2}{3}c$ weighed against the $\frac{1}{3}$ probability of losing the public good which is of value 1 to the agent. This results in a net change in expected utility from lying of $\frac{2}{3}c - \frac{1}{3}$. Thus, if $c > \frac{1}{2}$, then this decision rule is not incentive compatible.

In order to mitigate these incentive problems one has to adjust the probability that the project is undertaken when just one agent announces a valuation of 1. For the sake of illustration, take the cost to be $c = \frac{3}{4}$. The second best mechanism (achieving the maximal efficiency, while satisfying the incentive and participation constraints) is then as follows:

The Second Best Mechanism

		Agent 2's Value	
		1	0
Agent 1's	1	Build, Agent 1's Cost = $c/2$	Build Prob = $5/7$, Agent 1's Cost = c
Value	0	Build Prob = $5/7$, Agent 1's Cost = 0	Not Build

Next, let us consider what happens when the society is making decisions on two different projects at once.

Suppose that the agent's valuations are independent across the problems. If we do not link the decisions across the two problems, then the best we can do is just repeating the second best mechanism twice. For reference, this results in the following decisions as a function of the agents' types.

Two Public Decisions - UnLinked

Agent 2's values

		1	1	1	0	0	1	0	0	
Agent 1's values	1	1	1	1	1	5/7	5/7	1	5/7	5/7
		c/2	c/2	c/2	c	c	c/2	c	c	
	1	0	1	5/7	1	0	5/7	5/7	5/7	0
		c/2	0	c/2	-	c	0	c	-	
	0	1	5/7	1	5/7	5/7	0	1	0	5/7
	0	c/2	0	c	-	c/2	-	c		
0	0	5/7	5/7	5/7	0	0	5/7	0	0	
	0	0	0	-	-	0	-	-		

The top entries are the probabilities of undertaking the project

The bottom entries are agent 1's cost, conditional on building

Now let us try to improve on the above mechanism by linking the decisions.

First, let us examine how this works with a simple technique. Let us budget the agents so that they must act as if they have valuations that match the expected frequency distribution. We have to make an approximation, as there are only two problems and the empirical distribution is 2/3 on $v_i = 1$ and 1/3 on $v_i = 0$. As an approximation, we will require the agents to act as if they have a valuation of $v_i = 0$ on one of the problems and $v_i = 1$ on the other problem.

This results in the following linked mechanism. Here the outcomes are listed as a function of the true types, so this is the direct mechanism corresponding to one where each agent is treated as if they have a valuation of 1 on one of the problems and 0 on the other.⁷

A Linked Mechanism

⁷In this mechanism, we have coordinated the announcements in the most efficient manner. For instance, when a 0,0 meets a 0,1, we treat it as if the announcements are 0,1 and 0,1. This is incentive compatible.

Agent 2's values

		1 1	1 0	0 1	0 0
Agent 1's values	1 1	1 1 c 0	1 1 0 c	1 1 c 0	1 1 0 c
	1 0	1 1 c 0	1 0 c/2 -	1 1 c 0	1 0 c/2 -
	0 1	1 1 0 c	1 1 0 c	0 1 - c/2	0 1 - c/2
	0 0	1 1 c 0	1 0 c/2 -	0 1 - c/2	0 1 - c/2

This linked mechanism makes the efficient decision in all but one case. In the situation where both agents are of type 0,0, they are still treated as if they value at least one of the public goods, and so we mistakenly undertake one of the projects. This only occurs with probability 1/81. We have eliminated sixteen cases where the project was not always built when it should have been (and these are more likely cases).

The ex ante utility in this linked mechanism exceeds that of the unlinked case. The target total societal expected utility if the efficient decision were always taken is 4/3. The unlinked mechanism falls short of this by 4/63, while the linked mechanism falls short of this by 1/98.

We should be careful to point out the following issue. In the above mechanism, we end up violating the participation constraint (both interim and ex post) for the 0,0 type.

We can rectify this in several ways. One is simply to not undertake any projects when one of the two agents is of type 0,0. We can also simply go back to what the unlinked mechanism would do when at least one of the agents is of type 0,0. This results in the following mechanism, which improves over the unlinked mechanism, and yet satisfies both the incentive and participation constraints.

A linked mechanism satisfying the participation constraint

Agent 2's values

		1	1	1	0	0	1	0	0
Agent 1's values	1	1	1	1	1	1	1	5/7	5/7
		c	0	0	c	c	0	c	c
	1	0	1	1	1	0	1	1	5/7
		c	0	c/2	-	c	0	c	-
	0	1	1	1	1	1	0	1	0
	0	0	c	0	c	-	c/2	-	c
	0	0	5/7	5/7	5/7	0	0	5/7	0
		0	0	0	-	-	0	-	-

Issues of how to handle participation constraints in general will be discussed following the statement of the main theorem.

We next turn to a pure private good setting that, as shown by Myerson and Satterthwaite (1983), is also fundamental to demonstrating the tension between efficiency, incentive compatibility, and participation constraints.

EXAMPLE 3 *A Bargaining Problem*

This example is paradigmatic for bargaining (or bilateral monopoly) with uncertainty. A buyer and a seller must decide whether or not a good will be transferred from the seller to the buyer and what price will be paid by the buyer in the case of a transfer. There is uncertainty and the utilities are specified as follows: with probability $\frac{1}{2}$ the seller values the object at 0 and with probability $\frac{1}{2}$ she values the object at 8. With probability $\frac{1}{2}$ the buyer values the object at 10 and with probability $\frac{1}{2}$ he values the object at 2. Take these valuations to be independent.

Pareto efficiency dictates that the agents trade when the buyer values the object more than the seller, and not otherwise. To set a target, let us consider trade at a price that falls half-way between their valuations. This is captured in the following mechanism.

The Target (Ex Ante Efficient) Outcomes

		Buyer's Value	
		10	2
Seller's Value	0	Trade Prob = 1, Price = 5	Trade Prob = 1, Price = 1
	8	Trade Prob = 1, Price = 9	No Trade

However, this mechanism fails to be incentive compatible. It is easily checked that a seller with a valuation of 0 would prefer to be treated as if she had a valuation of 8. In fact, it follows from Myerson and Satterthwaite (1983) that there is no mechanism for this problem that is interim (or ex post) individually rational, incentive compatible, and Pareto efficient. The following mechanism adjusts the probability of trade to satisfy the incentive compatibility constraints, while still keeping with the equal splitting of the surplus from trade.⁸

An Incentive Compatible Mechanism

		Buyer's Value	
		10	2
Seller's	0	Trade Prob = 1, Price = 5	Trade Prob = $\frac{5}{8}$, Price = 1
Value	8	Trade Prob = $\frac{5}{8}$, Price = 9	No Trade

The efficiency loss of this mechanism is associated with the times when an 8 meets a 10 or 0 meets a 2 and exchange takes place with a probability below 1. Our goal is to show how the efficiency loss can be made to disappear when the buyer and seller bargain over many items where valuations are independent across items and agents.

As before, we link the decision problems. Require each agent to specify exactly $K/2$ objects for which he or she is “eager” to trade (corresponding to the valuations 0 for a seller and 10 for a buyer), and $K/2$ objects for which he or she is “reluctant” to trade (corresponding to the valuations 8 for a seller and 2 for a buyer); where we take K to be divisible by 2 for this example. Given these announcements, let us determine the outcomes by using the target outcome function on the announced valuations on each problem.

As before, there is an “approximately truthful” equilibrium of the linked mechanism, where agents tell the truth to the maximal extent that is possible given that their announcement must match the empirical frequency distribution while their realized type might differ from this distribution. The key insight is that, for instance, a seller could never gain by reversing announcements so that an 8 is announced as a 0 and a 0 as an 8. This will improve over the unlinked problem even when $K = 2$ (we leave the calculations to the reader this time), and again will approach full efficiency as K grows.

Let us add a further argument that *all* equilibria must lead to the same limit utility.

⁸This mechanism is not quite “second best”. If we adjust the prices to be 8 when an 8 meets a 10, and to be 2 when a 0 meets a 2, then we can increase the probability of trade up to $\frac{5}{6}$. That would offer an ex ante Pareto improvement over the mechanism that keeps prices at 1 and 9. Regardless of this change, the mechanism would still be inefficient, and there would still be improvements due to linking, just with some slight changes in the exact numbers.

Consider the seller. Suppose that the seller follows a strategy of announcing as truthfully as possible in the following way: if she has more than $K/2$ valuations of 0, then announce all of the valuations of 8 truthfully and randomly pick some surplus valuations of 0 to be announced as 8's; if she has fewer than $K/2$ valuations of 0, then announce all of the valuations of 0 truthfully and randomly pick some 8's to announce as 0's so as to meet the constraint. If she has exactly the right distribution, then announce truthfully.

Note that by using this strategy, *regardless* of what the buyer does, in the limit the seller will obtain her ex ante expected utility under the efficient mechanism on the problems where she is truthful, and her only potential losses (gains) are when she is forced to lie. That follows because the buyer must report something that agrees with the true expected distribution. If the seller is announcing as truthfully as possible in the manner described above, then the seller and buyer's announcements are independent. The seller has a strategy that guarantees her the ex ante efficient limiting payoff.

Given that by being approximately truthful, the seller can guarantee herself a given utility level, it follows that any strategy that the seller uses in any equilibrium must lead to at least this level of utility. Thus, any sequence of equilibrium strategies for the seller must lead to the same limiting payoff for her. By a similar argument the same is true for the buyer. Thus, each agent must get at least their ex ante expected payoff in any sequence of equilibria of the linking mechanisms. By the ex ante efficiency of these payoffs, it cannot be that either agent gets more. Thus *all* sequences of equilibria of the linking mechanism have the same ex ante limiting payoff.

Just as in the public goods example, if we hold our linking mechanism to satisfy the interim individual rationality constraint, then we need to make some modifications. An easy modification is to allow an agent to choose not to participate after learning his or her (overall) type. If one of the agents does not participate, then no trade takes place on any problems. This will happen with vanishing probability as K grows.⁹

3 A General Theorem on Linking Decisions

We now provide a theorem on linking decisions that show that efficiency gains can be made by linking any decision problems with any number of agents.

⁹The case of two point supports for both agents is somewhat degenerate, as for large K an alternative mechanism would simply be to allow for trade at a price of 5 on any item that either desired to trade, and then allow for participation decisions. However, with richer distributions of values such an alternative mechanism would no longer work. We discuss participation constraints in detail in what follows.

Let us first provide some definitions.

The Agents

Consider n agents who are involved in making decisions.

Decision Problems

A decision problem is a triple $\mathcal{D} = (D, U, P)$.

Here D is a finite set of possible alternative decisions; $U = U_1 \times \cdots \times U_n$ is a finite set of possible profiles of utility functions (u_1, \dots, u_n) , where $u_i : D \rightarrow \mathbb{R}$; and $P = (P_1, \dots, P_n)$ is a profile of probability distributions, where P_i is a distribution over U_i . We abuse notation and write $P(u)$ for the probability of u .

Note that the decision problem may involve many alternatives and the utility functions may have any possible structure. Thus, the coverage in terms of applications is quite general. The results extend easily to infinite settings through finite approximations.¹⁰

The u_i 's are drawn independently across agents. By considering this case, we can be sure that our efficiency results are not obtained by learning something about one agent's type through the reports of others (for instance, as in Crémer and McLean (1988)). We come back to discuss correlation.

Social Choice Functions

A *social choice function* on a social decision problem $\mathcal{D} = (D, U, P)$ is a function $f : U \rightarrow \Delta(D)$, where $\Delta(D)$ denotes the set of probability distributions on D .

This is interpreted as the target outcome function. We allow f 's to randomize over decisions since such randomizations admit tie-breaking rules that are common in the problems we have already discussed, among others.

Let $f_d(u)$ denote the probability of choosing $d \in D$, given the profile of utility functions $u \in U$.

Pareto Efficiency

A social choice function f on a decision problem $\mathcal{D} = (D, U, P)$ is *ex ante Pareto efficient* if there does not exist any social choice function f' on $\mathcal{D} = (D, U, P)$ such that

$$\sum_u P(u) \sum_d f'_d(u) u_i(d) \geq \sum_u P(u) \sum_d f_d(u) u_i(d)$$

¹⁰The extension is straightforward if outcomes lie in a finite dimensional Euclidean space, utility functions are continuous, and the social choice function is continuous in a suitable topology, the support of P is compact. Although it gets a bit tougher without continuity (or the compact support), presuming some measurability properties, the utility and social choice functions can still be approximated by simple functions, and things can be stated for arbitrarily large measures of utility profiles.

for all i with strict inequality for some i .

This standard property of ex ante Pareto efficiency implies the standard interim efficiency (conditional on each u_i) and ex post efficiency (conditional on each u) properties.

Linking Mechanisms

Given a base decision problem $\mathcal{D} = (D, U, P)$ and a number K of linkings, a *linking mechanism* (M, g) is a message space $M = M_1 \times \cdots \times M_n$ and an outcome function $g : M \rightarrow \Delta(D^K)$.

A linking mechanism is a mechanism that works on a set of decision problems all at once, making the decisions contingent on the preferences over all the decisions rather than handling each decision in isolation. Here M_i is a message space for agent i , and can be an arbitrary set. In the linking mechanisms we use to prove our results, M_i consists of announcements of utility functions for each decision problem, in fact a subspace of U_i^K .

We let $g_k(m)$ denote the marginal distribution under g onto the k -th decision, where $m \in M$ is the profile of messages selected by the agents.

Preferences over Linked Decisions

When we link K versions of a decision problem $\mathcal{D} = (D, U, P)$, an agent's utility over a set of decisions is simply the sum of utilities. So, the utility that agent i gets from decisions $(d^1, \dots, d^K) \in D^K$ given preferences $(u_i^1, \dots, u_i^K) \in U_i^K$ is given by $\sum_k u_i^k(d^k)$.

We assume that the randomness is independent across decision problems. Given independence and additive separability, there are absolutely no complementarities across the decision problems. The complete lack of interaction between problems guaranteed that any improvements in efficiency that we obtain by linking the are coming from being able to trade decisions off against each other to uncover intensities of preferences, and are not due to any correlation or complementarities.

Strategies and Equilibrium

A *strategy* for agent i in a linking mechanism (M, g) on K copies of a decision problem $\mathcal{D} = (D, U, P)$ is a mapping $\sigma_i^K : U_i^K \rightarrow \Delta(M_i)$.

We consider Bayesian equilibria of such mechanisms.¹¹

Approximating Efficient Decisions through Linking

Given a decision problem $\mathcal{D} = (D, U, P)$ and a social choice function f defined on \mathcal{D} , we say that a sequence of linking mechanisms defined on increasing numbers of linked problems,

¹¹We omit this standard definition. This and other omitted definitions may be found in Jackson (2003).

$\{(M^1, g^1); (M^2, g^2), \dots, (M^K, g^K), \dots\}$ and a corresponding sequence of Bayesian equilibria $\{\sigma^K\}$ approximate f if

$$\lim_K [\max_{k \leq K} \text{Prob} \{g_k^K(\sigma^K(u)) \neq f(u^k)\}] = 0.$$

Thus, a sequence of equilibria and linking mechanisms approximates a social choice function if for large enough linkings of the problems, on every problem the probability that the equilibrium outcome of the linking mechanism results in the same decision as the target social choice function is arbitrarily close to one. We emphasize that having the uniform convergence rate (so the maximum probability of a difference across problems is going to 0) is stronger than having what the linking mechanisms approximate the target social choice function in some average sense.

Strategies that Secure a Utility Level

An important aspect of our results is that *all* equilibria of our linking mechanisms will converge to being efficient. We establish this by showing that there exists a strategy that guarantees that an agent's payoff will be above a certain level, regardless of what strategy other players employ. We say that such a strategy secures a given utility level.¹²

More formally, consider an arbitrary mechanism (M, g) on K linked decision problems. A strategy $\sigma_i : U_i^K \rightarrow M_i$ secures a utility level \bar{u}_i if

$$E \left[\sum_{k \leq K} u_i(g^k(\sigma_i, \sigma_{-i})) \right] \geq K\bar{u}_i$$

for *all* strategies of the other agents σ_{-i} .

The Linking Mechanisms

We now give a description of the structure of the mechanism that will be used to approximate a given target social choice function f . The basic ideas behind the linking mechanism's structure have been outlined in the examples.

Consider K linked problems. Each agent announces utility functions for the K problems. So this is similar to a direct revelation mechanism. However, the agent's announcements across the K problems must match the expected frequency distribution. That is, the number of times that i can (and must) announce a given utility function u_i is $K \times P_i(u_i)$. With a finite set of problems, the frequency of announcements cannot exactly match P_i , unless $P_i(u_i)$

¹²This differs from the idea of a dominant strategy, and is closer in spirit to the concept of minimax strategy, although in a different setting.

divides K for each possible $u_i \in U_i$, and so we approximate P_i . The choice is then made according to f based on the announcements.

More formally, our K -th linking mechanism, (M^K, g^K) , is defined as follows.

Find any approximation P_i^K to P_i such that $P_i^K(v_i)$ is a multiple of $\frac{1}{K}$ for each $v_i \in U_i$, and the Euclidean distance between P_i^K and P_i (viewed as vectors) is minimized.

Agent i 's strategy set is

$$M_i^K = \{\hat{u}_i \in (U_i)^K \text{ s.t. } \#\{k : \hat{u}_i^k = v_i\} = P_i^K(u_i)K \text{ for each } v_i \in U_i\}.$$

Thus, agents must announce a vector of types across problems that matches the true underlying frequency distribution of their types.

The decision of g^K for the problem k is simply the target f operated over the announced types. That is, it is $g^K(m) = f(\hat{u}^k)$, where \hat{u}_i^k is i 's announced utility function for problem k under the realized announcement $m = \hat{u}$.

This is not quite the complete description of the mechanism. There are two modifications to the g^K 's that are needed. One is a slight change to realign the probability of choosing decisions to be as if announcements were exactly P_i rather than P_i^K , which is described in the proof. The other is an adjustment if there are at least three agents, that eliminates certain collusive equilibria and ensures that all equilibria converge to the desired targets. This is discussed after the theorem, and precise details appear in the proof.

Approximate Truth

The constraint of announcing a distribution of utility functions that approximates the true underlying distribution of types will sometimes force an agent to lie about their utility functions on some problems, since their realizations of utility functions across problems may not have a frequency that is precisely P_i . Nevertheless, strategies that are as truthful as possible subject to the constraints, turn out to be useful strategies for the agents to employ, and so we give such strategies a name.

We say that an agent follows a strategy that is *approximately truthful* if the agent's announcements always involve as few lies as possible. Formally, $\sigma_i : U_i^K \rightarrow M_i^K$ is *approximately truthful* if

$$\#\{k \mid \sigma_i^K(u_i^1, \dots, u_i^K) \neq u_i^k\} \leq \#\{k \mid m_i^k \neq u_i^k\}$$

for all $m_i \in M_i^K$ and all $(u_i^1, \dots, u_i^K) \in U_i^K$.

As we shall see, there always exists an equilibrium that involves such approximately truthful strategies. Moreover, such approximately truthful strategies are secure in that they guarantee the agent an average expected utility that converges to the ex ante efficient target

expected utility as K grows! This implies that all equilibria of the mechanism must converge to providing the same expected utility.

A Theorem on Approximating Efficient Decisions through Linking

Let $\bar{u}_i = E[u_i(f(u))]$, and let $\bar{u} = (\bar{u}_1, \dots, \bar{u}_n)$ denote the ex ante expected utility levels under the target social choice function. These are the targets for the utility level that we would like to implement.

THEOREM 1 *Consider a decision problem \mathcal{D} and an ex ante Pareto efficient social choice function f defined on it. There exists a sequence of linking mechanisms (M^K, g^K) on linked versions of the decision problem such that:*

- (1) *There exist a corresponding sequence of Bayesian equilibria that are approximately truthful.*
- (2) *The sequence of linking mechanisms together with these corresponding equilibria approximate f .*
- (3) *Any sequence of approximately truthful strategies for an agent i secures a sequence of utility levels that converge to the ex ante target level \bar{u}_i .*
- (4) *All sequences of Bayesian equilibria of the linking mechanisms result in expected utilities that converge to the ex ante efficient profile of target utilities of \bar{u} per problem.*


So, the theorem establishes that the linking mechanism has an equilibrium that is approximately truthful, that the outcomes of a sequence of such equilibria converge to the target ex ante efficient outcomes as the number of linked problems grows, that approximate truth secures the target expected utility level for any agent, and that all equilibria converge to provide these same utility levels.

The details of showing that there exists an approximately truthful equilibrium are the most involved in the proof (and we refer the reader to the proof for full details). However, this proof that all equilibria converge to the target expected utility levels is simpler and is argued as follows. Since the distribution of other players' announcements must match the underlying distribution of their types, and are independent from a given player's types, a given player knows that by announcing approximately truthfully, his or her expected utility matches that under complete truth, except to the extent that they are forced to lie. Thus, regardless of other players' strategies each player secures an expected utility level that is

approaching his or her target ex ante efficient expected utility level. Thus, any sequence of equilibrium strategies must give at least this secure level in the limit, as the player secure a level converging to this level simply by being approximately truthful.

This idea is pictured as follows, in the case of two players.

Figure 1.



approximate truth secures a level approaching \bar{u}_1

Let us emphasize that the availability of secure strategies makes the analysis quite robust. By being approximately truthful, an agent does not have to worry about what other agents do or even have any knowledge of what their potential types are or even what the distributions are. Beyond being focal, this makes it hard to imagine that any play besides approximate

truth could ever emerge. In fact, in the examples that have appeared important to us these are the only possible equilibria.

The fact that agents can secure payoffs with any approximately truthful strategy, also has interesting implications for the impossibility of improving through joint deviations. If each agent is playing an approximately truthful strategy, then the possible gain that might result from a joint deviation by some group of agents is bounded, as the remaining agents' utilities are secured regardless of the group's deviation. In fact, the structure of the mechanism that rules out collusion (described below) makes this true regardless of whether agents are playing approximately truthful or not. While this does not imply that every equilibrium is a strong equilibrium, it does imply that they are ε -strong equilibrium, in the sense that the gains from coalitional deviations approach 0 in the limit.

When there are three or more agents, there is a modification to the basic mechanism that is needed to ensure that all equilibria converge to the target limit. To see why we need such a modification, and what it should be, consider the following example.

EXAMPLE 4 *Collusion*

Let us reconsider the public goods example, Example 2. Recall that each agent has a value for the public good of $v_i = 1$ with probability $2/3$ and $v_i = 0$ with probability $1/3$; and that the cost of the public good is $3/4$. Suppose that we have linked K public good problems, where K is a multiple of 3.

As claimed in Theorem 1, there is an equilibrium where each agent is approximately truthful. If we simply run the mechanism as described above without any modifications, however, then for large enough n there also exist some other equilibria that do not converge to the target outcomes and expected utilities. Here is one such equilibrium. Let agents 2 through n always announce values of 1 on the first $2K/3$ problems, and values of 0 on the remaining problems. Have agent 1 announce truthfully on the last $K/3$ problems, and approximately truthfully on the remaining problems, randomizing on any necessary lies, but keeping any lies on the first $2K/3$ problems. It is easy to check that this is an equilibrium when n is at least 7.¹³

¹³Agent 1 is the only agent ever announcing positive values on the last $K/3$ problems. As such, agent 1 wants to be truthful on these problems as otherwise those public goods will not be built. Agent 1 is indifferent on the remaining announcements as those public goods will be built regardless. The other agents face the following choice. If they announce values of 1 on the first $2K/3$ problems, then they face costs of either $3/(4(n-1))$ or $3/(4n)$, each with probability $1/2$. If one of these agents deviated from the proposed equilibrium and announced a value of 1 on one the last $K/3$ problems instead, he or she would expect to pay a of $3/4$ with probability $1/3$ and a cost $3/8$ with probability $2/3$. This can result in a gain in utility in

This equilibrium results in only 8/9 of the problems being built, which is less than the efficient level when n is at least 7. Moreover, it leads to a very skewed distribution of costs in that agent 1 pays a disproportionate share.

So, how can we modify our basic linking mechanism to eliminate this and all other undesired equilibria in a simple way and without altering the mechanism’s nice efficiency properties? Here is such an approach. If we were running the mechanism and we saw such a sequence of announcements from agents 2 to n , we would think it highly likely that the agents had coordinated their strategies. What we can do is check agents’ announcements to see if they appear as if they match the joint distribution that would ensue under truth. If we find some agents whose joint announcements appear to be “too far from truth”, then we will simply ignore their announcements and randomly pick an announcement for them. We will occasionally make mistakes in doing this, but with a judicious choice of how to define “too far from truth”, we can keep the probability of this happening to a minimum and have this go to 0 in the limit. The full description of the modified mechanism appears in the proof in the appendix, and indeed it gets rid of all the undesired equilibria.

Note that the reason that such a modification is not needed with just two agents, is that with just two agents, the other agent’s announcements have to match the empirical frequency distribution, and then being approximately truthful secures a utility level converging to the efficient one. However, with more than two agents, we need to police joint announcements in order to ensure that the announcements that any given agent faces matches the empirical distribution. Once this is done, then approximate truth is again secures each agent’s ex ante efficient expected utility level in the limit.

4 Participation Constraints

Theorem 1 holds for any ex ante efficient social choice functions that we target. As such, f can satisfy any number of auxiliary properties, such as participation constraints¹⁴, fairness, etc.

The interest in participation constraints often arises in settings where agents have a choice

the first case as their the agent is pivotal as he or she changes the decision from not building that project to building it. However, in the second case, the agent does not alter the decision and ends up paying a higher cost than he or she would of if the value of 1 was announced on one of the first $2K/3$ problems. When n is at least 7, then it is a better response to announce a value of 1 on the first $2K/3$ problems.

¹⁴Participation constraints are also commonly referred to as individual rationality constraints. We use both terms.

of whether or not to participate in the mechanism, and this might occur after they already know their preferences. This, for instance, is often a constraint in any contracting setting, including the bargaining setting we considered in Example 3. If such participation decisions are relevant, then it is important to require that they be satisfied by our linking mechanisms all along the sequence, and not just in the limit.

Also, to make sure that the reason that our linking mechanisms are making improvements is due to the linking, and not due to relaxing a participation constraint, we want to show that we can strengthen Theorem 1 to hold when we impose various forms of participation constraints. This is important since in some settings the conflict between incentive compatibility and efficiency only arises when an interim participation constraint is in place, as was true in Examples 2 and 3.¹⁵

In order to be more explicit, let us provide formal definitions of the participation constraints.

Participation Constraints

Consider a decision problem (D, U, P) , where some decision $e \in D$ has a special designation, which may be thought of as a status-quo, an endowment, or an outside option. The interpretation is that e will be the default decision if some agent(s) choose not to participate. The standard formalization of this is via a participation constraint. Such constraints generally take one of three different forms, depending on the information that agents have available when they make their participation decisions.

A social choice function f satisfies an *ex ante participation constraint* (or an *ex ante individual rationality constraint*) if

$$E[u_i(f(u))] \geq E[u_i(e)]$$

for all i . Say that f satisfies a *strict ex ante participation constraint* if the above constraint holds strictly for all i .

A social choice function f satisfies an *interim participation constraint* (or an *interim individual rationality constraint*) if

$$E[u_i(f(u))|u_i] \geq u_i(e)$$

¹⁵This is true in a wide variety of collective decision problems, even with arbitrary transfers admitted. d'Aspremont and Gerard-Varet (1979) provide a key early theorem on this (see Jackson (2003) for a recent survey).

for all i and u_i . Say that f satisfies a strict interim participation constraint if in addition for each i there is at least one u_i such that the above constraint holds strictly.

A social choice function f satisfies an *ex post participation constraint* (or an ex post individual rationality constraint) if

$$u_i(f(u)) \geq u_i(e)$$

for all i and u . Say that f satisfies a strict ex post participation constraint if in addition for each i there is at least one u such that the above constraint holds strictly.

These constraints are increasingly demanding, with the ex ante constraint requiring that agents wish to participate in expectation, the second requiring that agents wish to participate after learning their own type, and the last one requiring that agents wish to participate in all circumstances.

Generally, the interim participation constraint is instrumental in the conflict between efficiency and incentive compatibility. Given that it is also often the sensible constraint in terms of modeling what agents know when they decide to participate in a mechanism, the interim constraint is the key one in much of the literature. We saw its role in Examples 2 and 3.¹⁶

Let us now discuss how our linking mechanism can be modified to satisfy any of these participation constraints.

Satisfying Participation Constraints with a Linked Mechanism

First, let us discuss why one needs some modification of the linking mechanism in order to satisfy a participation constraint. Reconsider Example 3. Suppose that we have linked 10 problems. So, a seller is required to announce five problems on which she has a value of 0 and five problems on which she has a value of 8. Suppose that the seller happens to be of a type that is all 8's. By participating in the mechanism (under any subsequent equilibrium play) with this type, she has a negative expected utility. Thus, in order to satisfy the interim participation constraint (or the ex post constraint),¹⁷ we need to modify the linking mechanism.

Consider a decision problem (D, U, P) with an option of not participating that results in a status quo option, denoted e . Consider the following variation on the mechanism (M^K, g^K) that is used in the proof of Theorem 1.

¹⁶The target outcomes we worked with there actually satisfied stronger ex post constraints, but in two-type models the two constraints are closely related.

¹⁷The ex ante constraint is satisfied.

In a first stage, the agents submit their actions from M_i^K , and decisions on all problems are given by $g^K(m^K)$. In a second stage, agents are each asked (say simultaneously) whether they wish to participate or not.¹⁸ If any agent chooses not to participate, then e is selected on all problems and otherwise the outcomes are $g^K(m^K)$.

We say that a strategy for an agent is approximately truthful, if m_i is approximately truthful and an agent chooses not to participate only in situations where his or her utility from non-participation (getting e on all problems) exceeds the utility of $g^K(m^K)$.

So, we have modified the linking mechanism to explicitly allow agents an option to not participate. We have done this at the ex post stage, which will provide for the strongest of the three forms of a participation constraint. It is important to note, however, that an agent must decide to participate in the whole linking mechanism or not to participate at all. We are not allowing an agent to pick some problems to participate in and not others. We return to discuss this shortly.

COROLLARY 1 *Consider any ex ante efficient f that satisfies a strict participation constraint of any sort: ex ante, interim or ex post. Consider the two-stage linking mechanisms with a participation decision as described above. For every K , there exists an approximately truthful equilibrium¹⁹ of the modified linking mechanism such that the resulting social choice function satisfies an ex post (and thus interim and ex ante) participation constraint, and the sequence of these equilibria approximate f .*

The proof of Corollary 1 follows fairly directly from the proof of Theorem 1, and so we simply outline it. Consider the approximately truthful equilibrium strategy identified in the proof of Theorem 1. Have agents play these strategies in the first stage of the mechanism. Next, let us describe the participation strategies for the second stage of the mechanism. Let m^K be the announcements from the first stage. If agent i 's utility²⁰ $\sum_k u_i^k(f(m^k))$ is at least $\sum_k u_i^k(e)$, then have the agent agree to participate; and have the agent choose not to participate otherwise. Given the “label-free” nature of the first stage strategies (see the proof of Theorem 1 for details), the choices of agents to participate are independent and do not affect the equilibrium structure of the first stage announcements. By a strong law of large numbers (e.g., the Glivenko-Cantelli Theorem), as K becomes large, the ex

¹⁸Whether agents are asked simultaneously or in sequence makes no difference for the conclusions of our result.

¹⁹The equilibrium notion is now Perfect Bayesian Equilibrium, as we are dealing with a two stage mechanism.

²⁰We are abusing notation slightly as f may be randomizing on outcomes, in which case we are looking at an expected utility.

post utility of approximate truth approaches the ex ante expected utility in the sense that $\frac{\sum_k u_i^k(f(u^k))}{KE[u_i(f(u))]}$ converges to 1 in probability. Also, $\frac{\sum_k u_i^k(e)}{KE[u_i(e)]}$ converges to 1 in probability. Given that some version of a strict participation constraint is satisfied by f , it follows that at an ex ante participation constraint is satisfied by f , and so $E[u_i(f(u))] > E[u_i(e)]$. Thus, $\sum_k u_i^k(f(u^k)) > \sum_k u_i^k(e)$ with probability approaching 1. This implies the claim.

Note that while Theorem 1 said that the profile of expected utility associated with all equilibria of the linking mechanism converged to the ex ante efficient utilities, there is no such claim in Theorem 1. With the added participation decisions of the agents, one introduces an additional equilibrium where all agents choose not to participate, anticipating that others will not participate. However, this involves a dominated strategy, as any undominated strategy requires participation whenever the utility from participation exceeds that of non-participation. Approximate truth coupled with a decision to participate whenever one has a higher utility from participating than the outside option, is still secure in the sense that it secures a utility level that approaches the ex ante efficient one provided other players participate. Thus approximate truth is still a good strategy, provided other agents don't follow dominated strategies.²¹

In the above analysis the participation decision is made at an ex post time, and hence we satisfy the strongest participation constraint. The analysis clearly extends to allowing this choice to be made at an interim or ex ante stage.

Difficulties with “Skimming”

The participation constraint(s) satisfied in Corollary 1 are of a form where an agent either participates in the linking mechanism as a whole, or does not. This is arguably the appropriate constraint. Nevertheless, one might wonder if we could satisfy an even stronger constraint, where we allow agents to participate on some problems and not others.

Suppose that instead of the above modification of the mechanism, we had instead allowed agents a choice of which problems they wish to participate in. One might conjecture, that the result would still go through. However, this is not the case, as seen in the following

²¹We remark that this does not allow one to conclude that all undominated (and hence trembling hand perfect) equilibria converge to the right utility levels. This can be seen in an example. Consider a decision problem where $D = \{a, b, c, d, e\}$, $U_1 = \{u_1, v_1\}$, $U_2 = \{u_2, v_2\}$, and the target f is given by $f(u_1, u_2) = a$, $f(v_1, v_2) = b$, $f(u_1, v_2) = c$, and $f(v_1, u_2) = d$. Normalize so that $u_1(e) = u_2(e) = v_1(e) = v_2(e) = 0$. Let $v_1(c) = v_2(d) = -10$, $u_1(c) = u_2(d) = 2$, and all other utilities be 1. Consider our linking mechanism with the participation decisions defined on $K = 10$ problems. There is an equilibrium where in the first stage agents say u_i on even problems and v_i on odd problems, and then choose to participate in the second period. What happens here is that if an agent deviates, they trigger a bad outcome that leads to non-participation in the second stage with a high probability. One can construct similar examples with larger K .

example.²² Again, reconsider the bargaining problem from Example 3, linked 10 times. Suppose that the buyer is announcing approximately truthfully, and then in participating on exactly those problems where the outcome is individually rational for them. Suppose the seller has a realization of 8's on the first five problems and 0's on the last five problems. Let her follow a strategy as follows. Announce 0's on the first five problems, 8's on the second five problems, and then not participate on the first five problems. The seller's explanation is that she has all 8's, and was simply forced to announce the 0's on the first five problems and would rather not participate on these problems. The outcome is that on the first five problems she keeps the goods (worth 8 to her), and on the last five problems ends up with an expected price of 4.5 (less those where the buyer does not participate), which is higher than the expected price of 3 (less those where the buyer does not participate) that she would have had if she had been truthful.²³

This shows that allowing agents to make their participation decisions on a problem-by-problem basis, rather than to participate in the mechanism overall, is problematic. This suggests that the linking mechanism needs to be treated as a single mechanism in terms of participation constraints, as otherwise agents can “skim.” We take this as an explanation for why agents are frequently not allowed to take part of a deal that they would like, such as some subset of a collection of mortgages that are bundled for sale.²⁴

5 Remarks and Discussion

A number of observations follow. These deal with an interpretation in which decisions are taken across time, implementing decisions as opposed to utilities, heterogeneous problems, rates of convergence, correlated types and problems, and informational requirements. These observations may be read selectively before turning to our related literature discussion in Section 6 and our conclusion.

Time

We have discussed our linking mechanisms as if all of the decisions were to be taken at the same time. When decisions are implemented across time, discounting is the effective

²²We thank Zachary Cohn for suggesting such an example.

²³We omit the detailed calculations with all participation decisions of the buyer incorporated, which show this to be an improving deviation.

²⁴It is a different problem to determine whether or not ex post constraints can be satisfied problem by problem, or whether they can all be satisfied with high probability. The latter question is the subject of Appendix 2.

determinant of the number of linked problems K .

To link K problems over time, consider the obvious variation of our previous mechanism made to operate over time. An agent is budgeted to announce a type of u_i exactly $P_i(u_i)K$ times. For example, consider our Example 1, with 12 periods and a vote in each period. Each agent is budgeted so they must announce $v_i = 2$ in exactly three periods, $v_i = 1$ in three periods, $v_i = -2$ in three periods, and $v_i = -1$ in three periods. Once an agent has used up his or her three announcements of a given type, he or she cannot announce that type in any of the remaining periods. A stream of decisions (d^1, \dots, d^K) results in a utility of $\sum_k \delta_i^k u_i^k(d^k)$ for agent i , where $\delta_i \in (0, 1]$ is a discount factor.

COROLLARY 2 *Consider a decision problem (D, U, P) and an ex ante efficient social choice function f with corresponding ex ante expected utility levels $(\bar{u}_1, \dots, \bar{u}_n)$. For any $\varepsilon > 0$ there exists \bar{K} such that for each $K \geq \bar{K}$ there exists $\bar{\delta}$ such that, for every $\delta \geq \bar{\delta}$, every Bayesian equilibrium²⁵ of the mechanism operating over time leads to an ex ante expected utility for each agent i that is above $\bar{u}_i - \varepsilon$.*

The fact that the equilibria of this mechanism operating over time have the desired properties follows as a straightforward application of the security part of Theorem 1.²⁶ Consider what happens if an agent follows an approximately truthful strategy. Here this means announcing the truth in each period, up to the point where an agent has exhausted his or her budget of a given type. In any period where the agent is forced to lie, have the agent pick one of his or her remaining types that (myopically) does best for him or her against the empirical distribution of the other agent's types. Again, *regardless* of what other agents might do, this secures an ex ante utility level that approaches the efficient one as K becomes large. Note also that this can be done in a quite myopic and simple-minded way by an agent. Honesty here really is a best policy.

²⁵This, of course, implies that the same is true of stronger solution concepts such as Sequential equilibrium or Perfect Bayesian equilibrium.

²⁶The handling of the collusion modification for $n \geq 3$ (see the proof of Theorem 1) is somewhat subtle, as the mechanism can only see the history of announcements up to the current time. With large enough K this can be handled as follows. At each time, examine the joint distribution through that time as per the test outlined in the proof of Theorem 1. The threshold of the test at any time k is $\nu^k > 0$. If some agents' announcements fail the test at some time k , then those agents' announcements will be made randomly by the mechanism (subject to the overall announcement constraints) until a time where the cumulative test is again passed. The thresholds of the test ν^k are set to converge to 0 in k , but also such that if agents are approximately truthful, then as K grows, then the expected frequency of periods in which the test will be failed goes to 0. By a theorem of Kolmogorov (13.4 in Billingsley (1968)), this can be done, for instance, by setting ν^k inversely proportional to the log of k .

Again, due to the security result, *all* equilibria must lead to utility levels that are at least as high as that from playing approximate truth, and so as K grows (or equivalently, letting the discount factor go to one) all equilibria must converge to providing the target utility levels. Note also, that this is true regardless of what agents know at what time. It is enough for agents to only know their preferences up to the current date to follow the secure strategy of approximate truth.

One might imagine that agents try to game the system, and, for instance, in the voting game from Example 1 use more of their 2's and -2's in early periods in the voting game than are realized. However, if there is not too much discounting, then the implications of the security levels are that such behavior will not be beneficial.

Let us emphasize here that the reasoning behind our results is quite different from that underlying repeated game folk-theorems. This can be seen in two ways. First, *all* of our equilibria approximate, in utility terms, a particular efficient target; while folk theorems generally lead to large sets of equilibria, some of which are very inefficient. Second, our results concern limits of finite repetitions in an arbitrary setting, rather than infinite horizon games.²⁷ This power derives from the fact that we are designing the mechanism, and can constrain the agents' announcements to match the target distribution; which is an entirely different point of view from that in the repeated game literature.

Outcomes and Utilities

While the theorem states that all equilibria lead to the same limiting utilities, and we know that the approximately truthful equilibria lead to the right limiting outcomes, we might want the even stronger conclusion that all equilibria lead to the same limiting efficient outcomes, not just the right limiting utilities. There are two things to say about this. First, and most important, for many problems, tying down the ex ante expected utilities to the target efficient utilities does tie down the outcomes as well. In fact, this is true for a generic choice of the utilities and probabilities (in the case where f is deterministic). Second, in cases where tying down the utilities does not tie down the outcomes, why might we care about the outcome when the utilities are already correct? The reason we might care would generally be that some unmodeled party has preferences over outcomes (for instance a cost of providing a good). If this is the case, then we can add that party to our setting and define the ex ante efficient rule accounting for their preferences too. Then applying our theorem will tie things down to ending up with the correct utilities for that agent too, and so outcomes

²⁷While there are some finite horizon folk-theorems (e.g., Benoit and Krishna (1985), they play off of multiplicity of equilibria in ways that bear no relationship to what we have analyzed here.

will be further determined to the extent necessary.

Heterogeneity in Problems

A clear limitation of the analysis that we have provided is that the decision problems being linked have the identical primitives. An easy generalization of the theorem is as follows. Consider a class of problems where the number of alternatives is bounded by some M . Now put a grid on such problems. For instance, let utility functions take on only some finite set of possible values and probabilities take on some finite set of possible values. In such a world, there is only a finite set of potential decision problems. Now, given any decision problem, match it with one in this class as closely as possible. Now, take an arbitrary set of K decision problems. As K becomes large, most of the decision problems will end up mapping to some problem which has many other problems matching to it too. So, all but a vanishing fraction of decision problems will have arbitrarily many other problems that look arbitrarily close to it. Now apply the theorem to each subsequence. We will get approximate efficiency on each subsequence. Thus, if decision problems all lie in some compact space, then we can take an arbitrarily heterogeneous sequence of them and the results of the theorem will apply to all but at most a vanishing fraction of them.

Moreover, one can see that we could get some partial improvements even in cases with limited numbers of repetitions and where the problems are all different, but have some relationship to each other. For instance, consider the case where there is a single seller who is bargaining with many different buyers. Each buyer is buying only one good, but the seller is selling many goods. Even though we cannot link the buyers' announcements, we can still link the seller's announcements to ensure approximate truth on her side. That will still lead to some improvements.

An open question for further research is the extent to which unrelated problems, such as a bargaining problem and a public goods problem, can be beneficially linked.

Large Numbers Reasoning

It is important to emphasize that the intuition behind the results here is quite distinct from other large numbers implementation theorems. That is, we know from the previous literature that increasing numbers of agents can, in certain circumstances, lead to increased competition and to efficient outcomes. Essentially the intuition there is that in the limit individual agents become negligible in terms of their impact on things like prices, so their incentives to try to manipulate the outcome to their advantage disappears.²⁸ In our linking of

²⁸See, for instance Roberts and Postlewaite (1973) and the literature that followed.

decisions the reasoning behind the gains in efficiency is quite different. Given that there is a fixed number of agents, they are not becoming negligible. In fact, they each hold substantial private information in terms of their overall ability to influence outcomes.²⁹ The key is that linking has helped us by giving a richer set of decision problems to trade-off against each other to help discover agents' preferences.

How Large is Large?

We can put a bound on the number of problems where any mistake will be made in the linking mechanism we have proposed here. The bound comes from what is known of laws of large numbers, such as a very useful theorem due to Kolmogorov.³⁰ Here it implies that the proportion of problems out of K on which agents might be forced to lie is of the order of $\frac{1}{\sqrt{K}}$. As we know that the secure strategies of approximate truth have lies that are then bounded by this, we obtain a crude upper bound on the distance from full optimality. It can be at most on the order of $\frac{1}{\sqrt{K}}$ in terms of percentage distance from full ex ante efficiency.

In many problems it is in fact closer. To get some feeling for this, let us consider a very simple example.

EXAMPLE 5 *Fast Gains*

Consider an object to be allocated to one of two agents. Each agent has a valuation for the object that is either 1 or 10, with equal probability, with independent draws. An ex ante efficient decision is to give the object to an agent with valuation of 10 if such an agent exists, and to either agent if both have valuations of 1. To be symmetric, flip a coin if both agents have a value of 10 or both have a value of 1. This results in a total ex ante expected utility across both agents of 7.75.

Without any linking, without transfers, and subject to incentive constraints, the best we can do is to flip a coin and randomly assign the object. This results in a total expected utility of 5.5.

We can also consider linking such decisions together. The following table provides the expected utility as a function of the number of the linked decisions.³¹

²⁹Thus, they are *not* informationally small in the sense of McLean and Postlewaite (2002).

³⁰See (13.4) in Billingsley (1968).

³¹The calculations here are for the “best” linking mechanism - one that minimizes the total number of misallocations subject to incentive constraints. In this example it is a variation on our previously described mechanism, where the mechanism helps in deciding where agents announce 10's if they have too few or too many compared to what they are allowed to announce. This actually corresponds to the choosing the best allocation subject to giving each agent half of the objects. Our previously described linking mechanism does

Number of Linked Problems:	1	2	4	6	limit
Expected Utility Per Problem:	5.50	7.19	7.50	7.69	7.75

Correlation

We have focussed on the case of independent types, both across agents and problems. The linking of decisions has helped even in the complete absence of any correlation. Thus, the intuition for why linking decisions together helps improve things has nothing to do with correlation in information being exploited. Nevertheless, it can still be that some forms of correlation make tradeoffs more or less likely, and thus more or less useful.

Correlation of any given agent’s preferences across problems (but not agents) does not impact the fact that approximate truth secures an expected utility level approaching the target level, provided the correlation is not perfect. However, correlation can help or hurt the speed of convergence. This is easily seen by looking at the extremes. If the problems are perfectly positively correlated, then there is no benefit to linking. Effectively, the all problems are exact copies of the *realization* of the first problem and so no tradeoffs across the problems are possible. So, it is clear that this is a worst-case scenario. On the other hand, if we examine the voting example, then perfect negative correlation across adjacent problems - at least in terms of intensities - is the opposite extreme and the best possible scenario. To see this, note that if we know that an agent cares intensely for one problem, then he or she will not care intensely for the next problem. Then we can ask an agent to declare which of the two problems they care more for, and there will be no difficulties at all - full efficiency can be attained. Looking at these two extremes suggests that there may be some sort of monotonicity in terms of the correlation structure. We work out the details showing that this is true for the voting example, Example 1, in an appendix.

In terms of correlation across agents, we already know from the mechanism design literature that having some correlation (across agents) can help in designing mechanisms, provided large rewards and penalties are possible and no ex post participation constraints are imposed (e.g., Crémer and McLean (1988)). The basic idea is that we can use one agent’s announcement to glean information about what the other agent(s) should be saying and thus to design incentives. In the absence of such rewards and punishments, our results suggest budgeting agents’ announcements to match the expected frequency distribution, and then checking the joint distribution of announcements to be sure that it does not stray excessively from the expected joint distribution. Indeed, approximate truth is easily seen to be a ε -Bayesian slightly worse than this one. We use the best linking mechanism only because it simplifies the calculations for this table, and with 6 linked decisions there are already 4096 utility profiles to worry about.

equilibrium of such a mechanism for any ε and large enough K . However, approximate truth may not be an exact equilibrium and no longer secures an expected utility level approaching the target levels, so that there may be quite varied equilibria.³²

What Does the Mechanism Need to Know?

As with all Bayesian mechanism design problems, there is a dependence of the mechanisms we suggest on the distribution of types, in this case the P_i 's. How robust are the mechanisms?

There are two things to say here. First, the security of approximately truthful strategies means that very little knowledge is required on the part of the agents. As long as the mechanism designer has performed approximately correct estimation of the P_i 's, an agent does not need to know anything about the possible types or distributions of the other agents, as approximate truth is a secure strategy.

Nonetheless, the mechanism itself still relies on the P_i 's. An important observation is that changing those P_i 's will change the secure payoffs in a *continuous* way. This means that mispecifications of the mechanism are not as problematic as with some other Bayesian mechanisms that are more precariously constructed. As the mechanism is increasingly misspecified, the expected utility levels will be further from the target efficient levels. Still, all equilibria will have to be above these secure levels.

To add to this, there are also many settings where the mechanism can be operated in ignorance of the distribution of utility levels. For instance, let us consider the voting example from before, and let us suppose that all we know is that agents' utilities fall in some interval $[-v, v]$, and have some symmetric distribution about 0 (so their intensities of preference for b matches that for a). Then we can operate a mechanism as follows. Have agents vote over the K problems. Next, let them also rank the problems 1 to K . Then, whenever there is a tie, break the tie in favor of the agent who has ranked the problem higher. This agent is likely to have had the more intense preference. Regardless of the distribution, this mechanism will have an approximately truthful equilibrium, and approximate truth secures the efficient level.

³²From the Bayesian implementation literature (in particular using the Bayesian monotonicity property; e.g., see Jackson (1990)), we can deduce that for some problems with correlation between agents types, any linking mechanism that has an equilibrium that approximates the target utility levels, will also have other equilibria that do not. Thus, there is a precise sense in which this is the best that we can do.

6 Related Mechanisms and Literature

One of the main conclusions from our results is that one should expect to see a linking of decisions problems in practice, as it can lead to substantial efficiency gains. It is thus not surprising that one can find, scattered about in the literature, examples of mechanisms that link problems. Some of these turn out to be cousins of the general mechanisms we have described here, but in the context of some particular problem.

Indeed, our initial discussions on this topic were spurred by trying to understand how the creative and innovative storable votes mechanism of Casella (2002) works. Casella's setting is one where a society makes binary decisions repeatedly over time, and in each period operates by a vote, choosing the alternative garnering the majority of votes. Her storable votes mechanism is one where an agent may store votes over time. So, an agent may choose not to vote in period 1, and then would have two votes at his or her disposal in period 2. Votes that are not cast by some agent in any period are stored for that agent's future use. Casella shows that in some cases, an equilibrium of the storable votes mechanism offers a Pareto improvement over separate votes. While there are many equilibria to this mechanism, and it is not clear that Pareto improvements are always present,³³ the idea behind the storable votes mechanism as a tool to get some gauge of the intensity of preferences, is what inspired our investigation into the linking of decisions.³⁴ Of course, the value added here is in showing how linking can be done to get full efficiency in the limit, how it can be done so that all equilibria converge, and that it applies to essentially any decision problem, not just a binary voting problem. In the context of the binary voting setting, our linking mechanisms also suggests some potential improvements relative to a storable votes mechanism. This comes from the way we force agents to ration their announcements. This might be thought of as giving voters a series of votes of different powers (corresponding to the relative intensities of their preferences). So, for instance, in Example 1, they would be given $K/2$ votes worth 2 points, and $K/2$ votes worth 1 point. They must spend exactly one vote on each problem. The key is that this additional rationing (eliminating the discretion of agents to freely pile up votes) leads to efficiency as an outcome and as the only viable outcome in the limit. More generally, depending on how variable preferences are, one could admit additional votes of

³³Experimental studies by Casella and Palfrey (2003) indicate that agents spend storable votes at least roughly in the right ways and realize some Pareto gains relative to standard voting mechanisms.

³⁴Hortala-Vallve (2003) (independently) studies what he calls "qualitative voting," which is a variation on Casella's storable votes that allows the transfer of votes freely across problems, whereas storable votes can only be stored for future use rather than borrowed from the future. As a result, in a two-intensity world some equilibria of his mechanism will Pareto dominate those of the storable votes mechanism.

various point levels, to match the distribution of intensities of preference.

Another setting where one sees mechanisms linking decisions across problems is one where a group of agents is splitting up a set of objects (as considered in Example 5). In such a context, under some symmetry assumptions, McAfee (1992) shows the limiting efficiency of a mechanism where agents take turns selecting objects.³⁵ With large numbers of objects, and symmetric type distributions across agents, this would lead to approximately the same outcomes as a linking mechanism that sought to give objects to agents with the highest valuation.

Let us close with some final remarks on the relation to some other literature that the linking of decisions might have brought to mind.

When thinking about voting problems and linking decisions, it is natural to think of logrolling.³⁶ Indeed there is some flavor of trading across decisions that is inherent in the linking mechanisms. However, logrolling generally has to do with some coalition (often a minimal majority) making trades in order to control votes, and usually at the expense of other agents. Logs are rolled in the context of majority voting mechanisms across different problems, which points out the important distinction that the mechanism itself is not designed with the linking in mind. This leads to a contrast between the benefits of linking mechanisms and the dark side of logrolling.

Finally, another place where some linking of decisions occurs is in the bundling of goods by a monopolist. The idea that a monopolist may gain by selling goods in bundles rather than in isolation was pointed out in the classic paper by Adams and Yellen (1976). Moreover, this gain can be realized when preferences over the goods are independent (see McAfee, McMillan and Whinston (1979)), can be enhanced by allowing for cheap talk where information about rankings of objects is communicated (see Chakraborty and Harbaugh (2003)), and in fact in some cases the monopolist can almost extract full surplus by bundling many goods (see Armstrong (1999)³⁷). Indeed, applying the linking decisions to the case of a bundling monopolist we can obtain (a strengthening of) Armstrong's result as a corollary to Theorem 1 by having the monopolist be agent 1 and the buyer be agent 2 and letting f

³⁵A related mechanism is discussed by Pesendorfer (2000) in the context of a set of bidders colluding in a sequence of auctions trying to decide who should win which auctions. See also, Blume (19), Campbell (1998), and Chakraborty, Gupta, and Harbaugh (2002), as well as earlier work on multimarket collusion by Bernheim and Whinston (1990).

³⁶For some of the classics on this subject, see Tullock (1970) and Wilson (1969), as well as the discussion in Miller (1977).

³⁷See also Fang and Norman (2003) who examine the efficiency gains from the bundling of excludable public goods.

be that the monopolist sells the good to the buyer at the buyer's reservation price whenever the reservation value is less than the cost of the good.

7 Conclusion

We have presented general theorems showing that the utility costs associated with incentive constraints vanish when a decision problem is linked with itself a large number of times. The argument is constructive in that we have demonstrated a mechanism that achieves these gains. The conclusions are strong in the sense that all equilibria of the canonical mechanism result in utilities that converge to the target utility vector. Indeed, we have argued that in a precise sense the canonical mechanism is almost preference revealing. We have made no claim that the canonical mechanism is the most efficient mechanism for achieving the utility gains associated with linkage, although a general rate of convergence is provided for it. Indeed, our examples demonstrate that one can frequently do better than the canonical mechanism. Several problems that appear to us quite interesting are left open. For example, are there simple canonical mechanisms that achieve rates of convergence similar to the most efficient linking mechanisms? If one cannot find such mechanisms in general, is there an interesting class of problems where the most efficient linking mechanisms take simple forms? Finally, we know relatively little that is general regarding what can be achieved when a small number of dissimilar problems are linked. Nevertheless, the general lesson is clear: whatever the utility costs associated with incentive constraints, the consideration of several problems at a time allows for trades that have agents reveal their private preferences, and as in the trade of normal commodities, this can help to improve the collective good.

References

- Adams, W.J. and J.L. Yellen (1976) "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics*, 90:3, 475-498.
- Armstrong, M. (1999) "Price Discrimination by a Many-Product Firm," *Review of Economic Studies*, 66:1, 151 -168.
- Benoit, J.-P. and V. Krishna (1985) "Finitely Repeated Games," *Econometrica*, 53, 905-922.
- Bernheim, B.D. and M.D. Whinston (1990) "Multimarket Contact and Collusive Behavior," *Rand Journal of Economics*, 21, 1-26.
- Billingsley, P. (1968) *Convergence of Probability Measures*, Wiley, New York, New York.
- Campbell, C.M. (1998) "Coordination in Auctions with Entry," *Journal of Economic Theory*, 82, 425-450.
- Casella, A. (2002) "Storable Votes," mimeo: Columbia University.
- Casella, A. and T.R. Palfrey (2003) "Storable Votes," mimeo: Columbia University and Caltech.
- Chakraborty, A., N. Gupta, and R. Harbaugh (2002) "Ordinal Cheap Talk in Common Value Auctions," mimeo: Baruch College CUNY, University of Michigan, and Claremont McKenna College.
- Chakraborty, A. and R. Harbaugh (2003) "Cheap-Talk Comparisons in Multi-Issue Bargaining," *Economics Letters*, 78, 357-363.
- Crémer, J. and R. McLean (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247-1257.
- d'Aspremont, C. and L.-A. Gerard-Varet (1979) "Incentives and Incomplete Information," *Journal of Public Economics*, 11, 25-45.
- Fang, H. and P. Norman (2003) "An Efficiency Rational for Bundling of Public Goods," Cowles Foundation Discussion Paper no 1441.

Hortala-Vallve, R. (2003) “Qualitative Voting,” mimeo: LSE.

Jackson, M.O. (2003) “Mechanism Theory,” in *Optimization and Operations Research*, edited by Ulrich Derigs, in the *Encyclopedia of Life Support Systems*, EOLSS Publishers: Oxford UK, [http://www.eolss.net].

May, K. (1952) “A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision,” *Econometrica*, 20, 680-684.

McAfee, R.P., (1992) “Amicable Divorce: Dissolving a Partnership with Simple Mechanisms,” *Journal of Economic Theory*, 56:2, 266-293.

McAfee, R.P., J. McMillan, and M.D. Whinston (1979) “Multiproduct Monopoly, Commodity Bundling and the Correlation of Values,” *Quarterly Journal of Economics*, 100:3, 371-383.

Miller, N.R. (1977) “Logrolling, Vote Trading, and the Paradox of Voting: A Game-Theoretical Overview,” *Public Choice*, 30, 51-75.

Pesendorfer, M. (2000) “A Study of Collusion in First-Price Auctions,” *Review of Economic Studies*, 67, 381-412.

Tullock, G. (1970) “A Simple Algebraic Logrolling Model,” *American Economic Review*, 60:3, 419-426.

Wilson, R. (1969) “An Axiomatic Model of Logrolling,” *American Economic Review*, 59, 331-341.

Appendix 1: Proofs

Proof of Theorem 1:

First, we offer a useful definition. Consider any K and the linking mechanism (M^K, g^K) . Let us say that a strategy $\sigma_i : U_i^K \rightarrow \Delta(M^K)$ for i is *label-free* if i 's strategy depends only on the realization of i 's preferences and not the labels of the problems. That is, if we permute which utility functions i has on each the problems, then we end up simply permuting i 's announcement in M^K in a corresponding manner.

Formally, given a permutation (bijection) $\pi : \{1, \dots, K\} \rightarrow \{1, \dots, K\}$ and any $u_i = (u_i^1, \dots, u_i^K) \in U_i^K$, let u_i^π be defined by $u_i^{\pi, k} = u_i^{\pi(k)}$ for each $k \in \{1, \dots, K\}$. So we have

just reshuffled the utility functions that i has under u_i on the different problems according to π . Given our definition of M_i^K there is a corresponding notion of m_i^π starting from any $m_i \in M_i^K$. A strategy σ_i for i is *label-free* if for any permutation $\pi : \{1, \dots, K\} \rightarrow \{1, \dots, K\}$ $\sigma_i(u_i^\pi)[m_i^\pi] = \sigma_i(u_i)[m_i]$, where $\sigma_i(u_i)[m_i]$ is the probability of playing m_i at u_i under σ_i .

The modification of the linking mechanism (M^K, g^K) for more than two agents is as follows.

For some j , $m_j \in M_j^K$ announced on linking mechanism (M^K, g^K) , and set of problems $T \subset \{1, \dots, K\}$, let $F_j^k(m_j, T)$ be the frequency distribution of announced types by j on problems in T . Thus, this is a distribution on U_j conditional on looking only at the announcements made on problems in T .

For any agent i and announced vector of $m \in M^K$ consider the following measure:

$$d_i^K(m) = \max_{j \neq i, u_j \in U_j, u_i \in U_i} |P_j^K(u_j) - F_j^k(u_j, \{k \in T | m_i^k = u_i\})|$$

If this measure is close to 0, then it means that the agents are not correlating their announcements. If it differs significantly from 0, then i is correlating announcements with some other agent j . That is, this measure looks at the conditional distribution of the announced u_j 's conditional on the dates that i announced some u_i and checks whether it is close to what the empirical distribution should be. It does this across all agents $j \neq i$ and all announcements of i .

By a strong law of large numbers of distribution, such as the Glivenko-Cantelli Theorem (see Billingsley (1968)), we can find $\varepsilon^K \rightarrow_K 0$, such that if agents are following strategies that are label-free, then

$$Prob \left[\max_i d_i^K(m_i) > \varepsilon^K \right] \rightarrow 0.$$

Modify the mechanism g^K as follows. For any i and announced $m \in M^K$ such that $d_i^K(m) > \varepsilon^K$, instead of using m_i , generate a random announcement \tilde{m}_i and then for each such i substitute \tilde{m}_i for m_i in determining the outcome. Do this by placing equal probability on each possible message in M_i^K , which results in an announcement that is independent of all other agents announcements and has the right frequency distribution.

We make one further modification of the mechanism. For a given K , the distribution P_i^K may not exactly match P_i . In order to make sure that for an arbitrary decision problem we always have an approximately truthful equilibrium, we need to be sure that the distributions exactly match P_i and not just approximately.³⁸ An easy modification of the linking

³⁸For some decision problems, this could turn out to make a difference, as it might be that f is ex ante efficient for the given P_i , but not for some approximations of it. This ex ante efficiency of f relative to an agent's expectations plays a role in obtaining an approximately truthful equilibrium.

mechanisms ensures this.

Find a smallest possible $\gamma^K \geq 0$ such that there exists another distribution \tilde{P}_i^K such that $(1 - \gamma^K)P_i^K + \gamma^K\tilde{P}_i^K = P_i$ (again noting that these can be written as vectors). Note that $\gamma^K \rightarrow 0$.

Now, on any given problem k let the mechanism g^K follow i 's announced m_i^k with probability $(1 - \gamma^K)$ and randomly draw an announcement to replace this with probability γ^K according to \tilde{P}_i^K , and do this independently across problems and agents. This means that the distribution of any i 's announcements that are used by the mechanism across problems will be *exactly* P_i .

Now, with a formal description of the mechanism in place, let us start by proving the second part of the theorem: that all sequences of equilibria converge to the same utilities.

Consider the following “approximately truthful” strategy σ_i^* . Consider a realized $u_i \in U_i^k$. For any $v_i \in U_i$ with frequency less than $P_i^k(v_i)$ in the vector u_i , announce truthfully on all problems k' such that $u_i^{k'} = v_i$. For other v_i 's, randomly pick $k \times P_i^k(v_i)$ of the problems k' such that $u_i^{k'} = v_i$ to announce truthfully on. On the remaining problems randomly pick announcements to satisfy the constraints imposed by P_i^k under M_i^k . By using σ_i^* agent guarantees him or herself an expected utility per problem approaching the utility that comes under truth-telling by all agents, regardless of the strategy of the other agents, as the agent is guaranteed that the distribution over other agents' types are approximately independently distributed and approximately what should be expected if the other agents were truthful (regardless of whether they are). Let \bar{u}_i be that utility level. As every agent can be obtain a limiting expected utility per problem of at least \bar{u}_i , regardless of the other agents strategies, by following the “approximately truthful” strategy σ_i^* , then it must that the lim inf of each agent's expected utility per problem along any sequence of equilibria is at least \bar{u}_i . However, notice that by ex ante efficiency of f , for any profile of strategies, and any k , if some i is expecting a utility higher than \bar{u}_i , then some j must be expecting a utility of less than \bar{u}_j . This implies that since the lim inf of each agent's expected utility for any sequence of equilibria is \bar{u}_i , it must also be that this is the limit of the expected utility of each agent, which is the desired conclusion.

To conclude the proof, let us show that there exists an “approximately truthful” Bayesian equilibrium of the linking mechanism such that the sequence of linking mechanisms and these corresponding equilibria approximate f .

Consider any agent i . If all agents $j \neq i$ play label-free strategies, then given the definition of the strategy spaces M_j and the independence across problems, the distribution of the

announcements of agents $j \neq i$ on any problem is given by P_{-i} , and this is i.i.d. across problems. Thus, for any best response that i has to label-free strategies of the other agents, there will in fact be a label-free best response for i .³⁹ Note also that any best response to some label-free strategies of other agents is a best response to *any* label-free strategies of the other agents. Given the finite nature of the game, for any set of label-free strategies of agents $-i$ there exists a best response for agent i , and, as argued above, one that is label-free. Thus there exists a label-free equilibrium.

Next, let us show that that there exists such an equilibrium that is approximately truthful in the sense that i never permutes the announcements of her true utility functions across some set of problems. Note that this together with the definition of M_i^k imply that as k becomes large the proportion of problems where i announces truthfully will approach one in probability. This again follows from distribution based versions of the strong law of large numbers such as the Glivenko-Cantelli Theorem, and will conclude proof of the theorem.

More formally, consider i 's label-free equilibrium strategy σ_i . Consider some $m_i = \hat{u}_i$ such that $\sigma_i(u_i)[m_i] > 0$. Suppose that there is some subset of problems $K \subset \{1, \dots, k\}$ such that i is permuting announcements on K . That is there exists a permutation $\pi : K \rightarrow K$ such that $\pi(k') \neq k'$ and $\hat{u}_i^{k'} = u_i^{\pi(k')}$ for all $k' \in K$. So i 's announcement under m_i reshuffles the true utility functions that i has under u_i on the problems K according to π .

Define \tilde{m}_i where this permutation on K is undone. That is, $\tilde{m}_i^{k'} = u_i^{k'}$ for each $k' \in K$ and $\tilde{m}_i^{k'} = m_i^{k'}$ for each $k' \notin K$. Then consider an alternative strategy (that will still be label-free) denoted $\tilde{\sigma}_i$ which differs from σ_i only at u_i and then sets $\tilde{\sigma}_i(u_i)[m_i] = 0$ and $\tilde{\sigma}_i(u_i)[\tilde{m}_i] = \sigma_i(u_i)[\tilde{m}_i] + \sigma_i(u_i)[m_i]$.

The claim is that $\tilde{\sigma}_i$ leads to at least as high an expected utility as σ_i . This follows from the ex ante efficiency of f . To see this note that the distribution of announcements under either strategy together with the strategies of the other agents is P on all problems and is independent across all problems (given the label-free nature of the strategies). Thus, the other agents' ex ante expected utilities on any given problem are not affected by the change in strategies. If i 's utility were to fall as a result of using $\tilde{\sigma}_i$ instead of σ_i , then it would be that f could be improved upon by a corresponding change of outcomes as a function of i 's utilities. This would contradict the ex ante efficiency of f .

Now we can continue to undo such permutations until we have reached a label-free strategy which has no such permutations. This is the ‘‘approximately truthful’’ strategy which we sought, and it still provides at least the utility of σ_i and thus is a best response, and since

³⁹Starting with any best response that is label dependent, any variation based on permuting the dependence on labels will also be a best response, as will a convex combination of such permutations which is label-free.

it is label-free it follows that the overall equilibrium is still preserved. Iterating on agents, leads to the desired strategy. ■

Appendix 2: Further Discussion of Participation Constraints

We can modify the mechanism so that not only is an ex post participation constraint satisfied overall, but also the probability that an interim constraint is satisfied on every problem at once by all agents, goes to one. This is done as follows.

Let $x(K)$ be such that $x(K)/K \rightarrow 0$, but

$$Prob[\exists v_i \in U_i : \#\{k | u_i^k = v_i\} > P_i(v_i)K + x(K)] \rightarrow 0$$

Let

$$\widetilde{M}_i^K = \{\widetilde{u}_i \in (U_i)^K \text{ s.t. } \#\{k : \widetilde{u}_i^k = v_i\} \leq P_i^K(v_i)K + x(K) \text{ for each } v_i \in U_i\}.$$

So, we have loosened the constraints on announcements to allow agents to over-announce some types, but at a vanishing rate.

The following lemma follows easily, given our earlier arguments.

LEMMA 1 *If f satisfies an ex post constraint, then when an agent i follows an approximately truthful strategy*

$$Prob[u_i(f(\widehat{u}^k) < u_i(e) \text{ for any } k] \rightarrow 0.$$

Thus, by being approximately truthful the probability that an agent will have an ex post individually rational outcome on every problem goes to one.

Lemma 1 does not claim that we will have an equilibrium that satisfies the ex post constraints if f is ex ante efficient. Indeed this is not generally true. To see why, note that since we have loosened the announcement constraints, in many settings agents will want to announce certain types on as many problems as permitted under the loosened constraint. So, even when their realized distribution is exactly the empirical frequency distribution, they will want to lie about some types to take advantage of the loosened constraint. In some settings, they will lie about types in a way that leads to an interim expected gain, but could possibly lead to an ex post bad outcome.

What is true, is that for generic problems⁴⁰, and for large enough K , if f is ex ante efficient then there will exist an equilibrium to the above described mechanism where the

⁴⁰Fixing D and the cardinality of each U_i , this is an open and dense set of Lebesgue measure one in terms of specifications of U_i 's, P_i 's, and f .

probability that every agent will have an interim individually rational outcome on every problem goes to one, and the proportion of problems where an ex post constraint fails goes to zero. The outline of the proof is as follows. Consider label-free strategies by the agents other than some i . The distribution of other agents' types is identical across problems and converges to the empirical one. For large enough K and a generic problem, the best responses of i are as if i faced the exact empirical distribution of the other agents on each problem. By the argument in Theorem 1, this implies that there are no problems on which agent i permutes his or her announcements of types. Identifying a label-free such best response (and one exists) provides the claimed equilibrium.

Appendix 3: Correlation

Consider a variation on the two-decision example presented above. First, let us draw agents' values on the first problem to be i.i.d. with equal probabilities on $\{-2, -1, 1, 2\}$. Next, we draw agent i 's value for the second problem, v_{i2} to be the same as for the first problem, v_{i1} , with probability $\rho \in [0, 1]$, and to be independent of the valuation for the first problem with probability $1 - \rho$.⁴¹

Now let us compare running separate voting mechanisms to running the linked mechanism where agents vote and also declare which problem they care more about or say that they are indifferent. Let us calculate the probability that a mistake is made under these two types of mechanisms. This is the probability that agents care in opposite directions on a given problem and with different intensities and a decision is made in favor of an agent who cares less about that problem.

Under separate voting mechanisms, the correlation pattern is irrelevant, and the chance that such an error occurs is $1/2$, conditional on agents caring in opposite directions and with different intensities. This situation arises $1/4$ of the time and so the total probability of such an error is $1/8$.

Under the linked mechanism, again the probability of this situation occurring is $1/4$. However, the chance that there is an error conditional on this situation arising is the $1/2$ times the probability (conditional on this situation) that the two agents have both announced

⁴¹This distribution is nicely symmetric and can also be described as picking the second problem valuations first and then drawing the first problem valuations in the manner described above.

“I care equally about the two problems”.⁴² The probability that this happens is

$$[\rho + \frac{1 - \rho}{2}]^2 = \frac{(1 + \rho)^2}{4}.$$

Thus, the overall probability of an error in this case is

$$\frac{(1 + \rho)^2}{32}.$$

When $\rho = 0$ this probability is minimized at $\frac{1}{32}$, and if $\rho = 1$ then this probability is maximized at $\frac{1}{8}$. Thus, the more positively correlated the valuations, the closer the linked mechanism is to just running separate mechanisms. The largest improvement comes from having independent values across the two problems.

This particular example does not allow for negative correlation, as things are either positively related or independent.

Let us consider another example where the correlation allows for a negative relationship between intensities.

The structure is parameterized by $\rho \in [-1, 1]$. Things are independent across agents. For a given agent i , we pick v_{i1} with equal probability on $\{-2, -1, 1, 2\}$. Next, we pick v_{i2} as follows. We first pick its sign. We do this in any manner so long as the marginal on positive and negative remains the same as the original distribution (equal probabilities). The correlation in signs will not matter in any way. Next, we pick the intensity of v_{i2} . We pick v_{i2} to have the same intensity as v_{i1} with probability $\frac{1+\rho}{2}$ and with the remaining probability of $\frac{1-\rho}{2}$ it is chosen to have a different intensity.

Here, it is easily seen that the probability of an error is

$$\frac{(1 + \rho)^2}{32}.$$

This is minimized at $\rho = -1$. So, negative correlation in intensities reduces errors to 0 and is even better than independence.

Appendix 4: Strategy-Proofness

We have shown that linking mechanisms can make improvements when we are discussing Bayesian incentive compatibility. As we now show, improvements are also possible when working with strategy-proofness (dominant strategy incentive compatibility).

⁴²Note that in this situation they will not have both named the same problem - they will either have named different problems or had at least one announce “equal”. The only potential error comes in when they both announced equality across problems.

A social choice function f is *strategy-proof* if $u_i(f(u_i, u_{-i})) \geq u_i(f(\hat{u}_i, u_{-i}))$ for all i , u_i , u_{-i} , and \hat{u}_i .

THEOREM 2 Consider two decision problems $\mathcal{D}^1, \mathcal{D}^2$ and corresponding strategy-proof social choice functions f^1, f^2 , where each $u_i^k \in U^k$ for each i and k is a strict preference over D^k . If $[f^1, f^2]$ is not ex post efficient viewed as a linked mechanism, then there exists a linked mechanism that Pareto dominates $[f^1, f^2]$ (from all time perspectives) and is strategy-proof.

We remark that theorem applies to $[f^1, f^2]$ which are ex post Pareto efficient when viewed separately, as long as they are not ex post efficient viewed as linked mechanism. So, for instance, this theorem applies to the voting problem of Example 1.

Proof of Theorem 2: Find some profile of utility functions u^1, u^2 and d^1, d^2 , where $[f^1(u^1), f^2(u^2)]$ is Pareto dominated by d^1, d^2 .

For any $1 > \varepsilon > 0$, define g^ε as follows. At any \hat{u}^1, \hat{u}^2 Let $g^\varepsilon(\hat{u}^1, \hat{u}^2)$ be a lottery with weight $(1 - \varepsilon)$ on $[f^1(\hat{u}^1), f^2(\hat{u}^2)]$ and ε on d^1, d^2 if d^1, d^2 Pareto dominates $[f^1(\hat{u}^1), f^2(\hat{u}^2)]$ at \hat{u}^1, \hat{u}^2 ; and let $g^\varepsilon(\hat{u}^1, \hat{u}^2)$ be $[f^1(\hat{u}^1), f^2(\hat{u}^2)]$ otherwise. It is clear from construction that g^ε strictly Pareto dominates f^1, f^2 from each time perspective. So, let us check that for small enough ε , g^ε is strategy-proof.

Consider some i and u_i^1, u_i^2 . If i lies and says $\tilde{u}_i^1, \tilde{u}_i^2$:

Case 1: $[f^1(u^1), f^2(u^2)] \neq [f^1(\tilde{u}_i^1, u_{-i}^1), f^2(\tilde{u}_i^2, u_{-i}^2)]$.

Here, by the strict preferences and strategy-proofness of f^1, f^2 , for small enough ε , there can be no gain in lying under g^ε .

Case 2: $[f^1(u^1), f^2(u^2)] = [f^1(\tilde{u}_i^1, u_{-i}^1), f^2(\tilde{u}_i^2, u_{-i}^2)]$.

Here, lying can only hurt, since the preferences of the other agents have not changed and the starting decisions from which g^ε is determined are the same, and so the change can only go against i 's preferences. ■