

The Linking of Collective Decisions and Efficiency

by Matthew O. Jackson and Hugo F. Sonnenschein *

Revision: March 28, 2003

Abstract

For groups that must make several decisions of similar form, we define a simple and general mechanism that is designed to promote social efficiency. The mechanism links the various decisions by forcing agents to budget their representations of preferences so that the frequency of preferences across problems conforms to the underlying distribution of preferences. We show that as the mechanism operates over a growing number decisions, the welfare costs of incentive constraints completely disappear. In addition, as the number of decisions being linked grows, a truthful strategy is increasingly successful and secures the efficient utility level for an agent.

*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 grant SES-9986190 is gratefully acknowledged, as is support from the University of Chicago where Jackson was visiting when this research began. We thank Bill Zame and the participants of SWET 2003 for helpful comments.

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 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 not commonly analyzed in terms of information and strategic play, and understood in terms of their ability to mediate incentive constraints.

One of the most fundamental lessons of this new understanding is that when classical problems, such as oligopoly, bargaining, and principal agent relationships, are repeated, the limits of cooperation commonly increase, and sometimes in the limit allow for outcomes that are fully efficient. This literature is built on an idea that is rooted in everyday experience: namely that failure of cooperation today can be answered tomorrow in a manner that encourages continual cooperation.

The purpose of this paper is to initiate a parallel investigation along a different dimension. We argue that similar gains in efficiency can be realized when classical incentive problems are linked, not temporally, but across separate problems or separate aspects of a social decision, exchange, or negotiation. We exploit the idea that when several independent problems are linked, or when there are several independent aspects of a given problem, then it makes sense to speak of rationing 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.

We should emphasize that while there is some parallel between linking decisions across problems and considering repeated games, the reasoning behind the results obtained here is very different from that underlying folk theorems. To understand how and why, it is useful to be more specific about our ideas and analysis.

Consider a set of decision problems, and suppose that agents' preferences are completely separable and independently distributed across these problems. This independence ensures that any improvements obtained in efficiency are not the result of some complementarities or correlation across problems, but really due to the method of linking itself. In our mechanism, agents are constrained to represent their types, which across a set of linked decision problems will be a vector, in ways that conform (as closely as possible) to the underlying distribution.

They 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. Within this constraint they play a Bayesian game across the linked problems, where they can misrepresent their preferences on individual problems but must respect the underlying distribution in the aggregate. Now, consider any fully (that is, ex ante) Pareto efficient social choice function that indicates what decisions we would like to make on each problem as a function of preferences, and leads to ex ante expected utilities of u_1, \dots, u_n for the agents. Generally, such an ideal rule will not be implementable in the presence of incentives on any individual problem. However, we show that this simple mechanism, based on the rationing of type announcements, has the property that *all* equilibria will approximate the desired function as the number of problems that are linked becomes large.¹ We do this by showing that truth is a very powerful strategy, as by being truthful (as closely as possible under the rationing of announcements) will secure an agent i his or her part of the ex ante expected utility, u_i .

In addition to the limiting results, we shall also see that significant efficiency gains are possible from linking even a few problems. We illustrate these small number gains in the context of a series of examples that will also make the approach and general results quite clear. Without further ado, let us turn to some examples and we will return to further discussion of our results and broader interpretations at the end of the paper.

2 Examples

Our first example is almost paradigmatic for voting theory and illustrates some of the main ideas.

EXAMPLE 1 *A Binary Decision - Voting Problem*

Consider two agents who are making a binary decision that affects both of their 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 decision is represented by $d \in \{a, b\}$.

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 problem is the difference in valuations

¹Note that this provides another important distinction from folk theorems repeated games. Rather than having many equilibria some of which are inefficient, here we have all equilibria converging to the same desired efficient utility levels.

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 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 more 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.

Let us start by considering a very simple and natural mechanism in this context. The two agents vote over the decisions a and b . If they both vote for the same decision, then that decision is taken. Otherwise a fair coin is flipped.

This mechanism has some nice features: it is incentive compatible (in fact dominant strategy incentive compatible) to vote for one's preferred decision. The mechanism is also anonymous and neutral. Moreover, it satisfies some efficiency criteria. It is an ex post Pareto efficient mechanism; that is, the decision that ends up being made is always Pareto efficient relative to the realized preferences. It also is a second best mechanism in that it maximizes the total sum of utilities, subject to the incentive constraints.²

It is important however to note that the mechanism is not ex ante efficient and does not maximize the total sum of agents utilities overall. 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; that is, has the intensity of 2. 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 and against the agent when the agent cares less.

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.

In terms of understanding the difficulties here, note the following feature. Regardless of whether an agent has preferences 1 or 2, he or she has the same preferences over any lotteries over the decisions - always preferring the lottery with more weight on decision b . There is no incentive compatible way to discover whether the agent cares a lot or a little about the

²It is the unique such mechanism which is also anonymous and neutral (a version of May's (1951) theorem). Other second best mechanisms are the dictatorial mechanism where one of the agents gets to pick the decision unilaterally (violating anonymity), and variations on the voting where a non-fair coin is used and, for instance, favors one of the decisions which is labeled the status-quo (violating neutrality).

decisions.

Linking Two Such Decisions

Next, let us consider a situation where there are two different 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 simply have considered a duplication of the previous decision problem.

One approach to solving this problem is simply 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. So one possible outcome of the coin flips results in a decision of (a_1, b_2) , which leads to a total utility of -2. This outcome is Pareto inefficient as both agents would prefer to have the decision of (b_1, a_2) , which leads to a total utility of 0.

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 allow each agent to announce only one utility of magnitude 2 (either a -2 or a 2) out of the two problems and require that the other utility be of a magnitude 1 (either a -1 or a 1). We then run the ex ante efficient mechanism on these constrained announcements. So, 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 utility has a larger magnitude and flip a coin in the event of a tie on magnitudes.³

³Note that we can also implement the above linked mechanism in the following manner. 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

It is straightforward to check that there is a Bayesian equilibrium 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 flips a coin and lies about the magnitude of the corresponding utility.

In fact, all equilibria 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.⁴

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 a total utility maximizing outcome. It turns out that 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 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. This 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;

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 a bargaining setting, insurance setting, and others.

⁴In 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 make comparisons after types are known, but before coins are flipped, then the linked mechanism dominates the separate voting.

- (c) agent 2 has equal magnitudes across the problems but not agent 1; and so this is analogous to sub-case (b);
- (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. Thus, linking the problems has cut the probability of making an error on each problem in half.

The reason why linking the two decisions together helps out is as follows. The inefficiencies in the separate problems are due to our inability to discover the agents' preference intensities. 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. Of course, if there are also ex post inefficiencies in the original problem, we can also try to fix those.

Let us make an important observation in this regard – no *interpersonal* comparability in utilities is needed in the above analysis. The ex ante inefficiency of the separate voting is not due to uncertainty regarding which of the two agents cares more - but rather 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 will 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 n independent decisions together of the type described above, where n 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 $\frac{n}{2}$ 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 a 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 $\frac{n}{2}$ problems, then the agent randomly picks $\frac{n}{2}$ 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 $\frac{n}{2}$ problems.

As n 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 the ex ante (and thus interim and ex post) efficient decisions.

As we shall argue below, this will in fact be true of all equilibria of this mechanism.

We should mention that the linking method we have proposed above can be further improved upon, by taking advantage of some specific aspects of the problem. Generally, we will not be proposing the best possible method of linking decisions, but we will propose a simple method that will reach full ex ante efficiency in the limit. To see that there are variations on this mechanism which perform slightly better along the sequence, but of course reach the same limit, consider the following improved (in fact optimal) version of a linking mechanism. 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 $\frac{n}{2}$ 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 $\frac{n}{2}$. 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 $\frac{n}{2}$ 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.

This voting example has provided some of the basic ideas that underlie more general results. To point out how this will work more generally in terms of setting, possible utility functions, numbers of agents, worrying about multiple equilibria, etc., we discuss some other examples before proving the general theorem.

EXAMPLE 2 *Taking Turns*

In the above example it is pretty clear that in equilibrium players wish to announce “approximately truthfully” under the linked mechanism, in that they will truthfully announce which problems they care about more intensely, except to the extent that they hit the constraint. Let us now show that this is not just an artifact of the two magnitudes of utility, but holds more generally. To keep the exposition simple, we consider a problem with three magnitudes of utility. We will see shortly that the results are fully general, and that will be pretty clear when seeing how the arguments work in this example.

Let us consider a situation with three agents dividing up an inheritance. There are a number of items, $m \geq 3$, to be divided among the agents. Agent i 's value for item j is denoted v_{ij} and takes on values in $\{1, 2, 3\}$, each with equal likelihood. The v_{ij} 's are independently distributed across agents and items.

If there were only one problem, then it is clear that there is no anonymous, incentive compatible, and ex ante efficient mechanism.⁵ Anonymity and ex ante efficiency require one to give the item to the agent who has the highest value for it, with some need for tie breaking. Clearly this is not incentive compatible as each agent would declare that he or she values the item at a level of 3. The best we can do respecting incentive compatibility and anonymity is simply to randomly assign the item. So here, if we try to operate things separately on each item, we end up simply randomly assigning items.

Let us link the decisions as follows. Require each agent declare a 3 for exactly $\frac{m}{3}$ items, a 2 for exactly $\frac{m}{3}$ items, and a 1 for the remaining items. Then operate the ex ante efficient and anonymous mechanism where each item is given to the agent with the highest valuation for that item with random tie breaking, based on these announcements.⁶ In this particular context, there is a very natural counterpart to this linking mechanism which is to randomly pick

⁵There is always an ex ante efficient and incentive compatible mechanism: always give all items to the first agent. So anonymity plays a role here. Alternatively, dropping anonymity note that there is no incentive compatible mechanism that maximizes the sum of the utilities across agents.

⁶A related idea would be to have agents submit rankings of objects. Chakraborty, Gupta and Harbaugh (2002) show how a mechanism based on providing rankings of objects can help a seller of multiple objects when trying to communicate values of those objects to prospective bidders in auctions. See also Chakraborty and Harbaugh (2003b), who explore benefits from rank orderings in cheap talk in sender receiver games.

an ordering over the agents and let them take turns in picking items - or in the terminology of sports: we hold a “draft”. That mechanism was studied by McAfee (1992) who showed that there is an equilibrium of the taking turns mechanism (which he called the alternating selection mechanism) which leads to efficiency in the limit. Indeed, these two methods are intuitively almost the same and lead to the same limiting distribution over the allocation of items in their “approximately” truthful equilibria. The advantage of the linking mechanism we describe, of course, is that while it is still simple it applies to any decision problem, well beyond the allocation of a set of indivisible goods.

In analyzing the mechanism(s) here, one can directly verify that there exists an “approximately truthful” equilibrium. This is not quite as obvious as it was in the previous example, and requires a bit of work. In particular, it would conceivably be advantageous for an agent to not announce a 3 on an item where he or she really has a value of 3 if he or she expected to get that item with high probability in any case, and then announce that 3 somewhere else to increase the probability of obtaining some other item. The full details of the argument appear later, but let us describe it loosely now. Suppose that the other agents are announcing approximately truthfully and randomly picking where to lie when they have to lie to meet the constraints on announcements. Then to a given agent the distribution over other agents’ announcements looks identical across problems. Given this, an agent cannot gain (and in fact would suffer in expectation) by permuting their true valuations; for example by reversing their valuations on two items, such as saying 3 when they have a 1, 2 when they have a 3, and 1 when they have a 2. They would end up trading probability of obtaining items that they value more for probability of obtaining items that they value less. This is easy to see here, and extends to other sorts of problems. In order to have some feeling for the key to that argument, note that it hinges on the fact that the decision that we are trying to implement is *ex ante* efficient and so in this problem is giving higher probabilities on items that are valued more highly. Having established that an agent does not want to permute announcements of his or her valuations, we have essentially shown that the agent wants to announce truthfully up to the constraint, and so we have a best reply of this form. There are a few details to be taken care of, but this is the essence of showing that there is a truthful equilibrium.

Again, as we increase the number of problems linked together we will converge to reaching full *ex ante* efficiency, as the proportion of problems where there are non-truthful announcements will go to 0.

While we have outlined why there exists an “approximately truthful” equilibrium of the linking mechanism, we might feel better if all equilibria of this linking mechanism must be

approximately truthful. In fact, we can show that *all equilibria* of this mechanism must lead to the same utility, and for this case that means they must all be approximately truthful. Rather than provide that argument in the context of this example, let us consider another example. There are two reasons for doing this. The first is that another example will further illustrate the breadth of coverage of our approach. The second is that the argument tying down all equilibria of the linking mechanism turns out to be different if there are two agents from when there are three or more. Let us start with the simpler intuition that underlies the two agent case. We return to discuss the uniqueness claim for three or more agent case in more detail after proving the main theorem.

EXAMPLE 3 *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. For simplicity, assume that each valuation occurs with equal probability and is independent across agents.

We would like to build the public good when $\sum_i v_i > c$ and not otherwise. Moreover, 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.

While our decision problem in terms of building the public project is a binary one, the decision in terms of allocating costs is more complex and so the number of outcomes is potentially quite large.

The desired decision rule that we have described will in general not be incentive compatible. To see this is quite straightforward. For instance, take the simple case where $n = 3$, $m = 1$ 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 $v_i = 1$. Consider an agent who has a valuation of $v_i = 1$. By pretending to have $v_i = 0$ that agent will still enjoy the public project with probability $\frac{3}{4}$, 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 both of the other agents have a valuation of 0, which happens with probability $\frac{1}{4}$. In particular the overall expected cost savings is $\frac{7}{12}c$ weighed against the $\frac{1}{4}$ 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 $\frac{7}{12}c - \frac{1}{4}$. Thus, if $c > \frac{3}{7}$, then this decision rule is not incentive compatible.

If the society is faced with making several such decisions, then we can link the decisions by requiring that agents announce values across the different problems that approximate the

frequency distribution. As the number of such linked decisions increases, we will converge to the first best efficient solution on each of them.

EXAMPLE 4 *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 $\frac{2}{3}$ the seller values the object at 0 and with probability $\frac{1}{3}$ she values the object at 8. With probability $\frac{2}{3}$ the buyer values the object at 10 and with probability $\frac{1}{3}$ he values the object at 2. Assume further that these valuations are independent.

It is fundamental since Myerson and Satterthwaite (1983) that there is no solution to this problem that is ex post individually rational, incentive compatible, and Pareto efficient. The following “second-best” mechanism maximizes the sum of the utilities subject to the constraints of individual rationality and incentive compatibility. When a 0-valued seller meets a 10-valued buyer exchange takes place at a price of 5. When an 8-valued seller meets a 10-valued buyer exchange takes place at a price of 8 only $\frac{5}{7}$ of the time; and $\frac{2}{7}$ of the time there is no exchange. Similarly, when a 0 meets a 2 exchange takes place at a price of 2 only $\frac{5}{7}$ of the time. If an 8 meets a 2, then there is no exchange. The efficiency loss of this mechanism is associated with the times when an 8 meets a 10 or 0 meets a 2 and no exchange takes place.

		Second Best Mechanism	
		Buyer's Value	
		10	2
Seller's Value	0	Trade Prob = 1, Price = 5	Trade Prob = $\frac{5}{7}$, Price = 2
8	Trade Prob = $\frac{5}{7}$, Price = 8	NoTrade	

Now, let us consider a buyer and seller who are bargaining over some number m of objects and in a situation where valuations are independent across items and agents. For simplicity, let us take m to be divisible by 3. In the same spirit as the previous examples, require each agent to specify the $2m/3$ times that he or she is “eager” to trade (corresponding to the valuations 0 or a 10), and the $m/3$ times that he or she is “reluctant” to trade.

But now, remove the probability of $\frac{5}{7}$ on the trades when eager and reluctant agents meet, so that trade happens with probability 1. So this is an ex ante efficient mechanism, provided the incentives are right for agents to announce their types approximately truthfully.

Ex Ante Efficient Mechanism
Buyer's Value

		10	2
Seller's	0	Trade Prob = 1, Price = 5	Trade Prob = 1, Price = 2
Value	8	Trade Prob = 1, Price = 8	NoTrade

Again, by the previous argument there is an “approximately truthful” equilibrium of the linked mechanism where agents are constrained to announce their valuations in proportion to the true distribution. And as argued before, we again have that this converges to full efficiency as the number of linked mechanisms increases.

But let us add a further argument that all equilibria must lead to the same limit utility. Consider the seller. Suppose that the seller follows a strategy of announcing approximately truthfully in the following way: if she has at least $2m/3$ 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 $2m/3$ 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 $2m/3$ constraint.

Note that by using this strategy, *regardless* of what the buyer does, in the limit the seller will obtain their full ex ante expected utility under the efficient mechanism. That follows because even if the buyer follows a strategy that depends on the labels of the problems, the buyer must report the correct distribution. If the seller is announcing approximately truthfully in the manner described above, then the seller and buyer's announcements are independent. Effectively, the seller has a strategy that guarantees her the ex ante efficient limiting payoff. 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 player 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.

Note that in the context of this example, as players are never indifferent, this argument actually also ties down the strategies in the limit to be approximately truthful.

Extending this uniqueness argument to more than two players requires an important but very natural modification of the linking mechanisms, as we discuss below.

We should point out that in the context of this example we have cheated a bit in showing that our linking mechanism leads to improvements. To be specific, we have not held our linking mechanism to satisfy the ex post individual rationality constraint except in the limit.

So while it does show that we can get closer to satisfying that constraint through linking mechanisms together, taking individual rationality seriously means that we might want to impose it on all of the linked problems as well. We shall come back to show that in fact we can reach the same conclusions even if we hold to the full ex post rationality constraint (that is even holding item by item) for any sized linking mechanism.

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.

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 .

The finiteness of the decision problems is assumed for ease of exposition as it provides for fairly clean and intuitive proofs. One way to extend the results to more general settings is directly through finite approximations.

We take the u_i 's to be drawn independently across agents. This makes achieving efficient decisions more difficult, as we know that correlation can help in designing incentive compatible and efficient mechanisms (for instance, as shown by Crémer and McLean (1988)). Thus, by considering cases with complete independence, we can be sure that our efficiency results are not obtained by learning something about one agent's type through the reports of others.

We abuse notation and write $P(u)$ for the probability of u .

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 $\Delta(\cdot)$ denotes the set of probability distributions on a given set.

We allow f 's to randomize over decisions since such randomizations admit tie-breaking rules that are natural in the problems we have already discussed, among others.

The notation $f(u)[d]$ denotes 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

$$\sum_u P(u) \sum_d f'(u)[d] u_i(d) \geq \sum_u P(u) \sum_d f(u)[d] u_i(d)$$

for all i with strict inequality for some i .

This is simply the standard definition of ex ante Pareto efficiency, and implies the standard interim (conditional on each u_i) and ex post versions (conditional on each u) as well.

Linking Mechanisms

Given a decision problem $\mathcal{D} = (D, U, P)$ and a number K of times that it is to be linked, a *linking mechanism* (M, \hat{f}) is a message space $M = M^1 \times \dots \times M^n$ and an outcome function $\hat{f} : M \rightarrow \Delta(D^K)$.

A linking mechanism is simply 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.

We let $\hat{f}_k(m)$ denote the marginal distribution under \hat{f} onto the k -th decision if the messages $m \in M$ are 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. Thus, there are no complementarities either in preferences or correlation across the decision problems. The complete lack of interaction between problems makes the gains from linking more difficult and really drives home the point that the efficiency gains we obtain are coming from being able to trade decisions off against each other to uncover intensities of preferences, and the gains are not due to any correlation or complementarities.

Strategies and Equilibrium

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

We consider Bayesian equilibria of such mechanisms.⁷

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, $\{(M^1, \hat{f}^1); (M^2, \hat{f}^2), \dots, (M^k, \hat{f}^k), \dots\}$, approximates f if there exists a corresponding sequence of Bayesian equilibrium $\{\sigma^k\}$ such that⁸

$$\lim_k \left[\max_{k' \leq k} E \left[\left| \hat{f}_{k'}^k(\sigma^k) - f(u^{k'}) \right| \right] \right] = 0.$$

Thus, a sequence of equilibria and linking mechanisms approximates a social choice function if for large enough linkings of the problems the equilibrium outcomes linking mechanism result in nearly the same decisions *on all problems* as the desired social choice function. We emphasize that being close on all problems is much stronger than having the average be close.

A Theorem on Approximating Efficient Decisions through Linking

We are now ready to present the main theorem. It is useful to first give a description of the mechanism that is used for the theorem. The basic ideas behind its structure have been outlined in the examples, and the linking mechanisms can be described as follows.

Each agent announces utility functions for the k problems. So this is like 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 approximately $k \times P_i(u_i)$.⁹ The choice is then made according to desired f based on the announcements.

The constraint of announcing a distribution of utility functions that approximates P_i will sometimes force an agent to lie about their utility functions on some problems, as just by chance their realizations of utility functions across problems may not match P_i . Nevertheless, the agent will still have strategies that are “approximately” truthful in a well-defined sense. To be precise, let us say that an agent follows a strategy that is *approximately truthful* if the agent's announcements are always such that they involve as few lies as possible. That is,

⁷We omit this standard definition.

⁸Note that mechanisms are distributions over finite set of decisions, and so distance between them is computed by viewing them as vectors.

⁹With a finite set of problems k , the frequency of announcements cannot exactly match P_i , unless $P_i(u_i)$ happens to be a fraction of k for each possible $u_i \in U_i$, and so we approximate P_i , as described in the appendix.

a strategy is approximately truthful if for any realization the number of problems on which the agent’s announced utility function and true utility function differ is minimized.

As we shall see, there always exists an equilibrium which involves such approximately truthful strategies. Moreover, such approximately truthful strategies are actually secure strategies in that they guarantee the agent the ex ante efficient expected utility! This implies that all equilibria of the mechanism must converge to providing the same expected utility.

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 on linked versions of the decision problem and corresponding (“approximately-truthful”) Bayesian equilibria that approximate f . Moreover, all sequences of Bayesian equilibria of the linking mechanisms converge to provide the same limiting expected utility per problem to all agents as they would obtain from truthful revelation under f on all problems. Furthermore, by following any approximately truthful strategy, an agent obtains his or her limiting ex ante efficient expected utility (as calculated under f) on each problem, regardless of the other agent’s strategies.*

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

EXAMPLE 5 *Eliminating Collusive Equilibria*

Consider the following three-person example.

Two decisions are possible, $D = \{a, b\}$. As in Example 1 we represent utilities in terms of the difference of utilities, $v_i = v_i(b) - v_i(a)$. Agent 1 always has $v_1 = -3$. So agent 1 is always in favor of decision a . Agents 2 and 3 are always in favor of decision b , but their utilities can each independently take on two possible values $U_2 = U_3 = \{1, 3\}$, each with probability $1/2$.

The solution we would like to implement in the utilitarian one where we choose decision a if $v_2 = v_3 = 1$ and otherwise we choose decision b .

Consider the linking mechanism over m linked versions of this problem, as we have described it in the previous examples. Agents 2 and 3 must each announce $\frac{m}{2}$ valuations of 1 and $\frac{m}{2}$ valuations of 3 over the m linked problems.

Agent 1’s announcement is always -3 for all problems, and so we can ignore it. As we have argued before, there is an approximately truthful equilibrium of the mechanism that results in our desired decisions with increasing probability as m becomes large.

However, note that in this example there is also another “collusive” equilibrium which does not result in our desired decision, and which involves coordination between agents 2 and 3. It is as follows. Have agent 2 announce $v_2 = 1$ on the even indexed problems and $v_2 = 3$ on the odd problems. Have agent 3 announce $v_3 = 3$ on the even problems and $v_3 = 1$ on the odd problems. This results in decision a being chosen on all problems! This is the best possible outcome for agents 2 and 3 and is clearly an equilibrium, but is not the utilitarian outcome that we desired.

So, how can we modify our basic linking mechanism to eliminate this bad equilibrium (and all other undesired ones) in a simple way and without altering its nice efficiency properties? Here is such a simple and natural approach. If we were running the mechanism and we saw a sequence of announcements from agents 2 and 3 where their total valuation turned out to be $3+1=4$ on so many of the problems, we would think it highly likely that this was not by accident but 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”, when 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, 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 under an ex ante efficient f , the mechanism results in a game that is essentially a strictly competitive one and so no collusion is possible.

4 Remarks and Discussion

Let us make a few remarks about the mechanism and the theorem’s coverage.

Rationalizability

The fact that any approximately truthful strategy secures an agent an expected utility that is approaching the ex ante efficient one, has some nice implications for the solvability of the game. We do not need to resort to Bayesian equilibrium or worry about player’s beliefs

about what strategies other players will play. The fact that they can secure a given payoff¹⁰ by following (any) approximately truthful strategy means that any rationalizable profile of strategies must lead to at least these secured payoffs.

Strong Equilibrium

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 players is bounded, as the remaining players' utilities are secured regardless of the group's deviation. In fact, the structure of the mechanism that rules out collusion makes this true regardless of whether players are playing approximately truthful or not. While this does not imply that any equilibrium is a strong equilibrium, it does imply that the gains from coalitional deviations will be limited and approaching 0 in the limit.

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 outcomes. There are two things to say on this. One is that for many problems, tying down the ex ante expected utilities does in fact tie down the outcomes. The other remark is that in cases where tying down the utilities does not tie down the outcomes, the reason we might care is that some other 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 and then apply the theorem.

Heterogeneity in Problems

Let us re-emphasize that the decision problems considered in Theorem 1 are completely arbitrary and so the coverage is quite general. This means that regardless of the nature of the problems and the reasons that efficiency might not be incentive compatible in isolated problems, linking the decisions together can improve. We have assumed that the decision problem being linked is the same in all cases. However, even if we have several different

¹⁰By secure we mean that the player gets at least that payoff regardless of the strategies of the other players. The game here is referring to the choice of σ_i 's by players as the strategy space and calculating payoffs from an expected utility point of view. Rationalizability is as defined by Bernheim (1984) and Pearce (1984).

problems, linking them is still advantageous, provided we obtain replications of each problem. It may be that one problem is a binary decision, while another is a bargaining problem, while another is a public good problem with transferable utility. Linking enough of these together will lead to efficiency gains.

Moreover, one can see that we could get some partial improvements even in cases 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.

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.¹¹ In our linking of 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.¹² 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.¹³ 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.

¹¹See, for instance Roberts and Postlewaite (1973) and the literature that followed.

¹²Thus, they are *not* informationally small in the sense of McLean and Postlewaite (2002).

¹³See (13.4) in Billingsley (1968).

In many problems it is in fact closer. To get some feeling for this, let us consider a very simple example. Consider an object to be allocated to one of two agents (a simplified version of Example 2). Each agent has a valuation for the object of either 1 or 10, with equal probability, with independent draws. An 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 an ex ante expected utility per agent of 3.875. Without any linking and subject to incentive constraints, the best we can do is to flip a coin and randomly assign the object. This results in an expected utility of 2.750.

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.¹⁴

Number of Linked Problems:	1	2	4	6	limit
Expected Utility Per Problem:	2.750	3.594	3.752	3.843	3.875

Other Desired Conditions: Individual Rationality

Due to the fact that the f 's can be any ex ante efficient mechanism that we desire, we can also satisfy whatever auxiliary properties we would like, such as individual rationality, fairness, etc. Moreover, in some cases we might want require that these conditions hold all along the sequence, and not just in the limit. This is quite natural for instance in the case of participation constraints such as individual rationality, and can be accommodated here.

As mentioned at the end of Example 4, the linking mechanism we described did not respect the individual rationality constraint except in the limit. Let us point out that in fact this is easily rectified. Let us run the linking mechanism, except for a change that we allow agents to walk away from any given problem if they do not like the outcome and then there is no trade (or some other status quo). This will guarantee that the outcome on every problem will be individually rational (from any time perspective). Agents will walk away on occasion under the linking mechanism, given that they are artificially constrained in the frequency distribution of their types and so even if f is individually rational the outcome might not always be so. However, which problems that the seller would like to walk away

¹⁴The 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 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.

from is not predictable by the buyer nor vice versa. This means that allowing players to veto the outcome of the mechanism to respect individual rationality, *makes no difference in the incentives across problems*. Thus, the claims about our linking mechanisms still hold even if we allow agents to walk away. In fact, as we link more problems, the fraction of problems where some agent walks goes to 0. Thus, all along the sequence we can also have individual rationality hold and still have our results be true.

It is very important to note that this logic extends more generally, and is not limited to the bargaining example. Individual rationality can be added to the main theorem itself in the same way. For instance in our public goods example, we can force no production any time some agent decides to walk away. This will occur in a vanishing fraction of the problems.

Correlation across Problems

When linking problems together what is the optimal mix of problems? Should we pick problems that are somehow related, or ones that are not? Thus far, we have focussed on the case of independent types. We know from the mechanism design literature that having some correlation *across agents* can often help in designing mechanisms, especially in situations where large rewards and penalties are possible and no ex post individual rationality constraints are imposed (e.g., Crémer and McLean (198?)). The idea is that we can use one agent's announcement to get some information about what the other agent should be saying and thus to design incentives.

Here, the linking of decisions has helped even in the complete absence of any correlation either across problems or across agents. Thus, the intuition for why linking decisions together helps improve things has nothing to do with correlation in information being exploited. As discussed above, the intuition instead stems from the ability to learn about intensities of preferences by exploiting tradeoffs across problems. Nevertheless, it can still be that some forms of correlation make tradeoffs more or less likely, and thus more or less useful. So, let us explore this in a bit more detail.

Let us first ask the question about correlation of each given agent's preferences across problems, while maintaining independence across agents. Two simple things are apparent and give us some idea of what we should expect. First, if the problems are perfectly positively correlated, then there is no benefit to linking. Effectively, the second problem is an exact copy of the *realization* of the first problem and so no tradeoffs across the two problems are possible. So, it is clear that this is a worst-case scenario. On the other hand, perfect negative correlation - at least in terms of intensities - is the opposite extreme and the best possible

scenario.¹⁵ To see this, note that if we know that an agent cares intensely for one problem, then they will not care intensely for the other 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. The following simple example shows this to be true.

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 first problem, v_{i1} , with probability $\rho \in [0, 1]$, and to be independent of the valuation for first problem with probability $1 - \rho$.¹⁶

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 "I care equally about the two problems".¹⁷ The probability that this happens is

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

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

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

¹⁵Perfect negative correlation in terms of intensities is a bit peculiar in terms of the overall distribution over values.

¹⁶This 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.

¹⁷Note 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.

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.

Some Comments on Related Mechanisms

In some cases, the linking mechanisms that we have defined take forms that have other interpretations or close cousins. For instance, in Example 2, the linking mechanisms results in a similar limiting distribution on outcomes as the taking turns mechanism studied by McAfee (1992). In the binary voting example, the linking mechanisms have features of a voting system where one had votes of varying power what could be spent on different problems, which is reminiscent of Casella's (2002) very innovative storable votes mechanism, although our linking mechanisms have some important distinctions in the way we force agents to ration their announcements to make sure that the equilibria are limiting efficient.

In other problems, such as the allocation of a private good, there are alternative mechanisms that operate in a very different way and can still attain efficiency, such as some auctions. With respect to that let us make two points. On the one hand our linking mechanisms operate without the need for any transfers or payments. This can be very important,

especially in situations where one is not sure to whom the revenue generated from the auction should be given. Auction revenue cannot be returned to the bidders without contamination of the incentives (or else some loss of individual rationality).¹⁸ On the other hand, the linking mechanisms do require a number of problems in order to get to full efficiency. Nevertheless, even for such private good allocation problems, the simple lessons from the linking mechanism can still be important: linking problems will enhance efficiency, and even simple budgeting of agents' actions or type announcements can help with incentives.

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.¹⁹ 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 is selling goods in bundles rather than in isolation is 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)). 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 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.

We wish to reiterate that our overall message goes beyond saying that linking decision

¹⁸As we know from d'Aspremont and Gerard-Varet (1973), there are some mechanisms that will be efficient and balanced (among participants, so in this case we can make that bidders), but we also know that such mechanisms will not satisfy even interim individual rationality constraints. One can see this fairly generally in Ledyard and Palfrey (2003).

¹⁹For some of the classics on this subject, see Tullock (196?) and Wilson (1969), as well as the discussion in Miller (1977).

problems can help enhance efficiency: it is also that the coverage of the linking mechanisms is broad, applying to most any setting; and that the ideas for discovering the preferences of agents on different problems by budgeting how they can act across different problems and imposing trade-offs can be useful and general tools for reconciling incentives with efficiency and other desiderata.

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. Nonetheless, the mechanism itself still relies on the P_i 's. Changing those P_i 's will generally change the secure payoffs in a continuous way, and so mispecifications of the mechanism are not as problematic as with some other Bayesian mechanisms that are more precariously constructed.

Even beyond this, we feel that the basic ideas here still provide some important insights into solving incentives problems. For instance, in the case of a series of binary decisions one can simply ask agents to rank order the problems in terms of the intensity of their preferences, and then use these rankings to help determine the outcomes. The important message is that the linking of decisions across problems offers the possibility of significant gains in efficiency. The realization of those potential gains might depend on the extent to which the decision making problem can be tailored to the environment.

References

- Adams, W.J. and J.L. Yellen (1976) “Commodity Bundling and the Burden of Monopoly,” *Quarterly Journal of Economics*, 90:3, 475-498.
- Mark Armstrong, M. (1999) “Price Discrimination by a Many-Product Firm,” *Review of Economic Studies*, 66:1, 151 -168.
- Bernheim, B.D. (1984) “Rationalizable Strategic Behavior,” *Econometrica*, 52: 1007–1028.
- Billingsley, P. (1968) *Convergence of Probability Measures*, Wiley, New York, New York.
- Casella, A. (2002) “Storable Votes,” mimeo: Columbia University.
- 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.
- Chakraborty, A. and R. Harbaugh (2003b) “Ordinal Cheap-Talk,” mimeo: Baruch College CUNY and Claremont McKenna College.
- Cr mer, 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.
- Ledyard, J. and T.R. Palfrey (2003) “A General Characterization of Interim Efficient Mechanisms for Independent Linear Environments,” mimeo: Caltech.
- 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.

Pearce, D. (1984) “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52, 1029-1050.

Tullock, G. “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

Proof of Theorem 1:

For any given k define the k -th linking mechanism, \hat{f}^k , as follows. This is the definition for $n = 2$, but we will state it for $n \geq 2$ as then the mechanism for $n \geq 3$ is an easily described variation.

For each i , P_i is the marginal distribution over the finite set U_i . Find any approximation P_i^k to P_i such that $P_i^k(u_i)$ is a multiple of $\frac{1}{k}$ for each $u_i \in U_i$, and the Euclidean distance between P_i^k and P_i (viewed as vectors) is minimized.

The mechanism \hat{f}^k is described as follows. Each agent i must announce u_i^k 's across different problems in a frequency exactly equal to P_i^k . Formally, i 's strategy set is

$$M_i^k = \{\hat{u}_i \in (U_i)^k \text{ s.t. } \#\{k' : \hat{u}_i^{k'} = u_i\} = P_i^k(u_i) \text{ for each } u_i \in U_i\}.$$

The choice of \hat{f}^k for the problem k' is $f^{k'}(\hat{u}^{k'})$, where $\hat{u}_i^{k'}$ is i 's announced utility function for problem k' under the realized announcement $m = \hat{u}$.

The modification of the mechanism for more than two players is as follows.

For some j , $m_j^k = \hat{u}_j \in M_j^k$ announced on linking mechanism \hat{f}^k , and set of dates $T \subset \{1, \dots, k\}$, let $\pi_j^k(\hat{u}_j, T)$ be the frequency distribution of announced types by j on dates in T . Thus, this is a distribution on U_j conditional on looking only at the announcements made on dates T .

For any k , agent i , and announced vector of \hat{u} and any k consider the following measure:

$$d_i^k(\hat{u}) \max_{j \neq i, u_i} |P_j^k - \pi_j^k(\hat{u}_j, \{t | \hat{u}_i = u_i\})|$$

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. 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 .

Let us say that a strategy σ_i for i is *label-free* if i 's strategy depends only on the profile of utility functions and not the labels of the problems. That is, if we permute which utility functions i has on the problems, then we end up simply permuting i 's strategy in a corresponding manner. Formally, Given a permutation (bijection) $\pi : \{1, \dots, k\} \rightarrow \{1, \dots, k\}$, 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 $m_i \in M_i^k$. Let us say that 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 .

By a strong law of large numbers of distribution, such as the Glivenko-Cantelli Theorem (see Billingsley (1968)), we can find $\varepsilon^k \rightarrow 0$, such that if agents are following strategies that are label-free, then the probability that $\max_i [d_i^k(\hat{u})] > \varepsilon^k$ goes to 0.

Modify the mechanism \hat{f}^k as follows. For any i and announced $m = \hat{u}$ such that $d_i^k(\hat{u}) > \varepsilon^k$, instead of using \hat{u}_i , generate a random vector \tilde{u}_i according to P_i^k and for each such i substitute \tilde{u}_i for \hat{u}_i in determining the outcome.

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 .

To do this, we need a 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 far enough along the sequence exactly match P_i and not just approximately.²⁰ Any easy modification of the linking mechanisms ensure this. Find a smallest possible γ^k 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 \hat{f}^k follow i 's announced $\hat{u}_i^{k'}$ 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, note that under this modification, all of our previous arguments still hold.

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 players, there will in fact be a label-free best response for i .²¹ Note also that any best response to some label-free strategies of other players is a best response to *any* label-free strategies of the other players. Given the finite nature of the game, for any set of label-free strategies of players $-i$ there exists a best response for player i , and, as argued above, one that is

²⁰For some decision problems, this could turn out to make a difference. The reason is that 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.

²¹Starting 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.

label-free. Thus there exists a label-free equilibrium.

Next, let us show 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 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 is still provides at least the utility of σ_i , so is still a best response, and since it is label-free it follows that the overall equilibrium is still preserved. Iterating on agents, leads to the desired strategy. ■

Strategy-Proofness

We have shown that linking mechanisms can make improvements when we are discussing Bayesian incentive compatibility - and in the proof of the limiting theorem we use a law of large numbers. As we now show, improvements are also possible when working with strategy-proofness (dominant strategy incentive compatibility).

[Insert definition of strategy-proof.]

THEOREM 2 Consider two decision problems $\mathcal{D}^1, \mathcal{D}^2$ and corresponding strategy-proof mechanisms 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 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.

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 \hat{f}^ε as follows. At any \hat{u}^1, \hat{u}^2 Let $\hat{f}^\varepsilon(\hat{u}^1, \hat{u}^2)$ 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 $\hat{f}^\varepsilon(\hat{u}^1, \hat{u}^2)$ be $[f^1(\hat{u}^1), f^2(\hat{u}^2)]$ otherwise. It is clear from construction that \hat{f}^ε strictly Pareto dominates f^1, f^2 from each time perspective. So, let us check that for small enough ε , \hat{f}^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 \hat{f}^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 \hat{f}^ε is determined are the same, and so the change can only go against i 's preferences. ■