

Efficient Teamwork

Endre Csóka

Alfréd Rényi Institute of Mathematics, Budapest, Hungary

Abstract

Our goal is to solve both problems of adverse selection and moral hazard for multi-agent projects. In our model, each selected agent can work according to his private “capability tree” involving hidden actions, hidden chance events and hidden costs but providing contractible consequences which are affecting each other’s working process and the outcome of the project. We will construct a mechanism that induces truthful revelation of the agents’ capabilities and chance events and to follow the instructions about their hidden decisions. This enables the planner to select the optimal subset of agents and obtain the efficient execution, which requires that each hidden decision of each agent should genuinely depend on the hidden capabilities and earlier hidden chance events of all other players. We will construct another mechanism that is collusion-resistant but implements an only approximately efficient outcome. The latter mechanism is widely applicable, and the major application details will be elaborated.

Keywords: Robust mechanism design, Stochastic dynamic mechanisms, Tendering systems. *JEL:* C72, C73.

1 Introduction

Assume that we want to manage a complex project by hiring some agents for different tasks. The agents have separate but interdependent working processes. For example, they might have to share common resources (e.g. machines or loading areas). Or a subtask of an agent must precede another subtask of another agent. Here an efficient cooperation means a stochastic decision plan which chooses each decision considering the current state of the entire project.

However, the current state of the project is not observable, because the capabilities, the chance events and the decisions of each agent are hidden from all other players. For example, if an agent finishes late (or with any unfavorable outcome), then we cannot ascertain whether he reported better capabilities or higher efforts than the truth or just had unfortunate chance events. Similarly, if an agent is able to finish a subtask earlier for a little extra cost, or he can adjust his usage of a common resource to the changing demands of others, then he can deny these capabilities or report them more costly.

We are offering a solution to resolve all these cooperation failures, namely, we design a mechanism that incentivizes the agents to truthfully reveal all their detailed capabilities and all chance events, and to follow the instructions of the planner. In the main part of the paper, we will show it on an idealized model, but in the appendix, we will show how it can be used under more realistic circumstances: if the players are not risk-neutral, they have non-quasi-linear utilities, they have limited liability, they do not precisely know (or are not able to define) their own capabilities, or we do not have unlimited computational capacity.

In our model, there is a principal and some agents. The principal can decide which agents to involve in the game, the others leave the game with utility 0. Each of the remaining players play a dynamic stochastic game according to their capabilities, which exert contractible influences




Figure 1: The capability tree of an agent describing the potential executions of his working process.

on each other. The payment to each agent is a function of these contractible events and the communication history.

A capability tree is a dynamic stochastic process describing the entire working capabilities of one agent. This contains decision nodes where the agent should choose the branch to continue, and there are chance nodes where the branch is chosen randomly with given probabilities. Each decision node and chance node has an absolute point in time.

A simple example for a capability tree is described by the decision tree in Figure 1, as follows. At time point 1, the agent has to choose between two possible working processes. The left process has a total cost of 3; and with probabilities 0.6 and 0.4, the process finishes with a success or failure, respectively. The right process costs 4, but the probabilities are 0.7 and 0.3 for the success or failure, respectively. In this simple example, the agent knows only the prior probabilities of the results until time point 5 (left process) or 7 (right process), when he gets to know the result. This capability tree is importantly different from a same tree with different time points, as we will see in Figure 2.

The capability tree, the decisions, the chance events and the cost of the work are all private information of the agent. In other words, from the point of view of everybody else, each agent looks like a black box which can communicate and outputs results, but about whom nothing else is observable. For example, the agent could easily choose the first, cheaper process instead of the second one (as instructed), or he could even report a completely different capability tree without the risk of being caught.

This was a simplified description of the capability trees. Here, the only consequence provided by the agent was a binary result of his work. Consequences in general, including direct influences between the capability trees of different players will be introduced in Section 4.2.

2 Example

We show a simplified model of the entire project. There is a central player called the principal, and some competing agents. Each agent privately knows his his capability tree. Contractible communication between the players is available throughout the game. The principal is free to choose which agents to involve in the project, and the others leave the game. Then all chosen agents execute their capability trees in parallel, each of them provides a result. At the end, the mechanism determines signed transfers t_i from the principal to each agent, as a function of the results and the entire history of communication. The utility of each agent is the transfer t_i he gets minus his costs. The utility of the principal is a joint valuation function of the results of the agents minus all transfers to them.

As a very simple example, consider the following two-task project. The principal must choose one agent per task. Each task i has binary results: success (\checkmark_i) or failure (\times_i). If both tasks succeed, then the principal gets $v(\checkmark_1\checkmark_2) = 20$. But if either fails, then the result of the other task is irrelevant, and the principal receives $v(\checkmark_1\times_2) = v(\times_1\checkmark_2) = v(\times_1\times_2) = 0$.




Figure 2: The capability trees of the agents and the evaluation of the truthfully reported trees.

Agents *A* and *B* apply for the first task, Agents *C* and *D* apply for the second task. Their capability trees are described in Figure 2. In words, Agent *A* would complete the first task with a cost of 5. Agent *B* would complete the first task with a cost of 6, but he would have an option of quitting at time point 2 (absolute time point, say, the end of the year 2002), with 0 cost. Agent *C* would have a cost of 1, until time point 3 (end of 2003), when he either gets to know that he failed to complete the task successfully (result X_2), or he will successfully complete it for a further cost of 2 (result \checkmark_2 with a total cost of 3), each option happens with probability 1/2. Agent *D* has the same capability tree but with 1 higher costs, and he observes his result at time point 1 (end of 2001).

We show the execution of the game under our revelation mechanism when all agents report their capability trees truthfully, and they act truthfully and obediently throughout the game. At the beginning, each agent is asked to report his capability tree. If the principal were to choose Agents *A* and *C*, then for an expected total cost of $5 + 2 = 7$, the principal would get $v(\checkmark_1\checkmark_2) = 20$ with probability 1/2. This would provide $1/2 \cdot 20 - 7 = 3$ expected total utility. If the principal were to choose Agents *A* and *D*, then the expected total cost would be $5 + 3 = 8$, and the expected total utility would be $1/2 \cdot 20 - 8 = 2$. If the principal were to choose Agents *B* and *C*, then the efficient strategy would be that Agent *B* should choose the left option, with a total cost of $6 + 2 = 8$, and the expected total utility would be $1/2 \cdot 20 - 8 = 2$. However, if the principal were to choose Agents *B* and *D*, then the efficient strategy would be for Agent *B* to choose the left option if Agent *D* succeeds, or the right option if Agent *D* fails. This way, the expected total cost would be $(6 + 0)/2 + 3 = 6$, and the expected total utility would be $1/2 \cdot 20 - 6 = 4$.

In general, the expected total utility can be calculated by a simple recursion as follows. The “combined capability trees” in Figure 2 show the joint execution of the capability trees of all pairs of agents for the two tasks. Such a tree describes all possible joint executions of the two capability trees. For example, the path to the third leaf of the joint capability tree of Agents *B* and *C* describes the execution that Agent *B* chooses the second option (completing the task) at time point 2, and then Agent *C* has the left outcome (failure) of the chance event at time point 3. The leaf has the pair of the corresponding results: success of the first task and failure of the second task; and we have a total cost of $6 + 1 = 7$. The total utility is the value of the pair of results (20 if both succeed, 0 otherwise) minus the total cost. For example, at this leaf, the total utility would be $v(\checkmark_1X_2) - (6 + 1) = 0 - 7 = -7$; for the next (rightmost) leaf, this is $v(\checkmark_1\checkmark_2) - (6 + 3) = 20 - 9 = 11$.

The state of the project can be represented by a point in the graph of the combined capability

tree. The value of a state of the project (indicated in small font brown integers) is defined to be the maximum expected total utility assuming an efficient strategy profile. Alternatively, it can be defined recursively from the bottom to the top, as follows.

- The value of a leaf is calculated as the joint value of the results minus the sum of the costs.
- The values of two points lying on the same edge of the decision tree (with no decision or chance node between them) are the same.
- The value at or just before a decision node is equal to the maximum of the values at the edges connected to the decision node from below.
- The value of a chance node is equal to the mean of the values of the edges connected to the chance node from below, weighted by their probabilities.

The value of the starting state is called the joint value of the pair (or subset, in general) of capability trees.

Once the values have been calculated, the project is executed in the following way. The principal chooses the pair (subset) of agents with the highest value and considers their joint reported capability tree. At each decision node, the principal observes which of the branches has the highest value, and asks the corresponding agent to choose that branch. At each chance node, the corresponding agent reports the outcome of the chance event, and the principal pays him the change in the value of the current reported state of the project, or in other words, the signed difference between the values of the states of the project after and before the chance event. At the end, each agent has to provide the reported result (otherwise he is caught for cheating and has to pay a huge penalty; or see Appendix E.1), and the principal pays him the reported cost at the leaf. This is what we call the first-price revelation mechanism.

We define the second-price revelation mechanism as the same except that each agent gets an additional second-price compensation: the difference between the value of the chosen subset minus the highest value of the subsets excluding the agent. In our example, the values of pairs are: $V(A, C) = 3$, $V(A, D) = 2$, $V(B, C) = 2$ and $V(B, D) = 4$, therefore, the principal chooses the pair (B, D) . The best pair without B is (A, C) , the best pair without D is also (A, C) , therefore, both agents get second-price compensations $V(B, D) - V(A, C) = 1$.

Consider the execution under the second-price revelation mechanism. We assume that the agents choose truthful strategies which will be shown to be a quasi-dominant equilibrium. According to the evaluation, agents B and D are asked to execute their capability trees. First, Agent D reports the outcome of his chance event at time point 1. If he succeeds, then he will get $10 - 4 = 6$ for this chance event, and Agent B will be asked to choose the first branch. But if Agent D fails, then he gets $(-2) - 4 = -6$ (he pays 6) for this chance event, and Agent B will be asked to choose the second branch. In total, Agent B receives his reported cost plus his second-price compensation, which is $0 + 1 = 1$ or $6 + 1 = 7$ payment depending on whether he was asked to choose the first or second branch, respectively. Agent D receives $4 + 6 + 1 = 11$ or $2 - 6 + 1 = -3$, namely, he gets 11 or pays 3 depending on whether he succeeds or fails at his task, respectively. The valuation of the result for the principal is 20 or 0, and she pays $7 + 11 = 18$ or $1 + (-3) = -2$, so her utility is $20 - 18 = 0 - (-2) = 2$ anyway.

With the first-price revelation mechanism, we expect the agents choosing the truthful strategies except that they report slightly higher costs, or in other words, they require slightly higher payment. For example, Figure 3 shows an example for the reported capability trees, called **proposals**, and the evaluation of them. We assume that each agent reports an additional cost 1. We call them **fair proposals with profit 1**.




Figure 3: The fair proposals of the agents with profit 1 and the evaluation under the first-price revelation mechanism.

In our example, this means that we have the very same process except that Agents *B* and *D* get 1 more payment for the higher reported cost, but they do not get the second-price compensations.

In this example, if one of these four agents reported his capability tree truthfully, then he would not be able to cheat later. We emphasize that this is not the case in general, at all. See for example Figure 1, where the agent can choose the (cheaper) left option instead of the (more expensive) right option, without the risk of being caught. This issue did not appear in our example just because we wanted to keep the example as simple as possible. We also note that these high risks are only the properties of our simple and extreme example. Most real-life projects contain a large number of smaller risks, providing much smaller total risk. We can think about the example so that we have hundreds of small independent identical projects of this kind. For further discussion of this topic, see Appendix E.4.

2.1 The offer-based mechanisms

We define the first-price offer-based mechanism as follows. The agents send offers for contract to the principal. The principal signs some of these contracts, and the other agents leave the game. The principal commits in advance to choose the strategy that maximizes her utility in the worst possible case. Under this mechanism, the proposals in Figure 3 translate to the following offers for contract.

Proposal of A: “ I will provide the result \checkmark_1 and the principal should pay me 6. ”

Proposal of B: “ If the principal sends me a message “start” until time point 2, then I will provide the result \checkmark_1 and she should pay me 7. Else I will provide the result \times_1 and she should pay me 1. ”

Proposal of C: “ The principal should send me a message before time point 3 about an amount of money x . At time point 3, I will send a message “yes” or “no”. If I send “yes”, then I will provide the result \checkmark_2 and the principal should pay me $4 + x$. If I send “no”, then I will provide the result \times_2 and the principal should pay me $2 - x$. ”

Proposal of D: “ The principal should send me a message before time point 1 about an amount of money x . At time point 1, I will send a message “yes” or “no”.

If I send “yes”, then I will provide the result \checkmark_2 and the principal should pay me $5 + x$. If I send “no”, then I will provide the result \times_2 and the principal should pay me $3 - x$. ”

(Negative payments mean positive payment in the opposite direction).

The execution is the following. The principal accepts B and D and sends “6” to D . At time point 1, D sends “yes” or “no” to the principal. In the case of “yes”, the principal sends “start” to B , otherwise she sends no message (meaning “do not start”). According to the contracts, the same payments are applied as with the revelation mechanisms, and the utility of 2 is guaranteed for the principal, this is her maximin utility.

The second-price offer-based mechanism means the same as the first-price one except that the principal should pay an extra amount of money to each agent, which money is equal to the difference between her maximin utility with all proposals and her maximin utility with all but the agent’s proposal.

Under both mechanisms, each agent promises a result in each case meaning that he accepts a huge penalty in the case if he provides a different result. We emphasize that this does not mean at all that the agent is forced to tell the truth. He can cheat about costs, probabilities. But even if he tells the truth in the beginning, he might be able to cheat about his decisions. E.g. in the case of Figure 1, even if the principal asks for the right branch, the agent is free to choose the left branch without the risk of being caught. This is very clear from his truthful proposal, as follows.

Truthful proposal with the capability tree in Figure 1:

“ The principal should send me a message “Left” or “Right” before time point 1.

Case Left. If she chooses “Left”, then she should choose an amount of money x before time point 5 (and send a message to me about it). At time point 5, I will send a message “yes” or “no”. If I send “yes”, then I will provide the result \checkmark and the principal should pay me $3 + 0.4 \cdot x$. If I send “no”, then I will provide the result \times and the principal should pay me $3 - 0.6 \cdot x$.

Case Right. If she chooses “Right”, then she should choose an amount of money x before time point 7. At time point 7, I will send a message “yes” or “no”. If I send “yes”, then I will provide the result \checkmark and the principal should pay me $4 + 0.3 \cdot x$. If I send “no”, then I will provide the result \times and the principal should pay me $4 - 0.7 \cdot x$. ”

Every proposal in the revelation mechanisms corresponds to an equivalent proposal in the offer-based mechanism, but it is not true in the other way. For example, in the offer-based mechanism, a risk-averse agent D , instead of asking for $5 + x$ or $3 - x$ payments in the two cases, he can ask for $5 + x$, and $3 - c + \frac{c^2}{c+x}$ with some (large) constant c . Such offers gently push the execution plan towards a less risky plan, or a plan where more risks are taken by less risk-averse agents.

However, under idealized assumptions including risk-neutrality, our results under the revelation and the offer-based mechanisms are the same, and most of the proofs are analogous. Still, the revelation mechanisms are easier to handle, therefore, we will prove our results with revelation mechanisms, and we will sketch the differences in Appendix C.

2.2 Incentive-compatibility under the second-price mechanism

We show here why the truthful strategy profile is a Nash-equilibrium under the second-price revelation mechanism. This will be the core idea of the proof that the truthful strategy profile

is also a quasi-dominant equilibrium, which is a new solution concept, a slightly weaker version of dominant strategy equilibrium.

Assume that, say, agent D changes his strategy, and the other agents play truthfully. For calculation purposes, let us ignore the second-price compensations to the other agents, but keep it to D . Now the expected (modified) utility of each agent is 0, because they get their costs paid, plus for all of their chance events, they get 0 more in expectation. If D did not participate, then the (modified) utility of the principal would be 3 by accepting A and C . If D makes a winning offer, and the value of the best pair of proposals increases to $3 + x$, then the principal will pay a second-price compensation x to D , and therefore, the (modified) utility of the principal will remain 3.

This means that the deviation of D from his truthful strategy does not change the expected (modified) utilities of all other players. Therefore, the expected loss of efficiency caused by his deviation must be equal to the expected loss in the utility of D compared to his truthful strategy.

2.3 Another example

We sketch another example which shows much more about the power of our results.

Assume that we have k mines of different goods. We need g_i number of goods from mine i , and our valuation $v(t) = C - \alpha \cdot t$ linearly decreases by the finishing time t when we get all (the last one) of the goods. There are miners applying for the mining jobs, and each mine has a fixed capacity of miners (maybe 1 or ∞). Goods are discrete units and mining is a Poisson-process: the probability distribution of future success of finding goods depends only on the abilities of the miner about that mine and on his efforts, but independent from the previous efforts and successes. Efforts are measured by cost-per-time rates.

The abilities and the effort levels of the miners are hidden, all what we can observe are the goods they find. Moreover, we assume only that we will observe the total amount of goods found by each miner, but the miners can report earlier or later finding times.

Our goal is to choose and allocate the miners and choose their effort levels optimally, namely, to maximize the expected valuation of the result minus the expected total costs of the miners.

Our mechanisms applied to this problem are the following. At the beginning, the miners report their abilities. Then we calculate the best plan assuming that the reports were truthful, and we allocate the miners accordingly. We ask the miners for the effort level in the plan (but we cannot observe whether they follow it) and to report whenever a good is found. Each time a good is reported to be found the miner gets an amount of money for the success. But there is a deduction rate for the time spent.

All the suggested effort levels, the prices for finding each kind of good and the deduction rates are recalculated whenever any of the miners report that he finds a good. These can be calculated as follows. With a simple backward induction, we calculate the expected total utility $v(\mathbf{g}')$ for each demand vector $\mathbf{g}' = (g'_1, g'_2, \dots, g'_k)$ assuming that the abilities of the miners are the same as what they reported, and we have full control on their actions. Then the price to be paid to the miner for finding a good of type i is $v((r_i - 1, \mathbf{r}_i)) - v(\mathbf{r})$, where r_i is the remaining demand of the good of type i . The penalty rate for the time spent is $\lambda_i(f) \left(v((r_i - 1, \mathbf{r}_{-i})) - v(\mathbf{r}) \right) - f$, where f is the suggested effort level (cost-per-time rate) and $\lambda_i(f)$ is the reported intensity of finding goods at effort level f . Beyond these payments, under the second-price revelation mechanism, the miners get second-price compensations for their initial proposals. Or under the first-price revelation mechanism, the fair offer includes a demand for a constant additional payment.

We note that if we had tried to solve the problem directly, without using any of the results in this paper, then it would have seemed to be a reasonable idea to also consider the tool of

one-time compensations to the agents whenever the price for finding a good or the deduction rate changes. But our results show that efficiency can be implemented (in quasi-dominant equilibrium) without using this tool. Without our result, it would be difficult to explain why we do not need this tool.

3 Comparison with the related literature

We consider the following related questions.

Problem 0. *Our model for single-agent projects with fixed result: auction problem.*

Consider the case when the principal has to choose a single agent to complete the entire project, and the only possible result is the completion of the task. This problem does nothing about cooperation failures or moral hazard, this is just a single-item auction problem for a negative-valued good, which good is the commitment for completing the task.

The second-price single-item auction (or Vickrey-auction) implements the efficient outcome in dominant strategy equilibrium: every agent makes a bid simultaneously, and the agent with lowest bid wins the task for the payment equal to the second lowest bid. [7]

If we want to maximize revenue for the principal, namely, we want to get the task completed for the lowest possible expected price, then it is a more difficult problem. In the Bayesian game, the optimal prior-dependent mechanism was found by Myerson [5], but it is used only as a benchmark for more realistic mechanisms. One of the most commonly used mechanisms is the first-price single-item auction, which is simply that all agents bid with a price simultaneously, and whoever bids with the lowest price wins the task for that price. The first-price single-item auction does not implement the efficient allocation of the good, and calculating the equilibrium strategy profile and expected revenue is a difficult task even for moderately simple prior distributions. The theoretical background why this is one of the most common mechanisms in practice is not entirely complete, but this is believed to be a reasonably good solution.

First- and second-price auctions for Problems 0, A and B are special cases of our first- and second-price mechanisms.

Problem A. *Our model for single-agent projects: selling the project to an agent.*

Consider the case when the principal has to choose a single agent to complete the entire project, but the agent works according to his dynamic stochastic capability tree with different possible results (e.g. completion time). Here, there are no potential cooperation failures, but moral hazard can be a problem. For example, the agent may be free to choose a different effort level than the socially efficient one (or the first branch instead of the second one in Figure 2), because none of his actions, costs, luck (chance events) and capabilities are observable, only his result.

The valuation of the principal on the possible results (consequences) $c \in \mathcal{C}$ is denoted by $v: \mathcal{C} \rightarrow \mathbb{R}$. The following class of mechanisms \mathcal{M} solves the moral hazard problem. We use an arbitrary single-item selling mechanism where the good to sell is the contract “if the agent provides a result c , then the principal will pay him $v(c)$ ”. For example, using a first-price single-item auction, the mechanism is the following. Each agent submits a bid b and the agent with the highest bid wins the task and will get $v(c) - b$ payment. These mechanisms guarantee that whoever wins will be incentivized for choosing his hidden decisions (e.g. effort levels) in the socially efficient way.

This means that considering mechanisms from \mathcal{M} reduces Problem A to Problem 0. For example, the second-price auction version implements efficiency in dominant strategies. However, it does not mean that \mathcal{M} always includes a best mechanism for maximizing the expected utility of the principal. It is not hard to construct a (not very natural) prior distribution and a

mechanism which beats all mechanism in \mathcal{M} for that prior distribution. Accordingly, it is very hard to say anything about how good the first-price auction mechanism is compared to other mechanisms, but this is believed to be reasonably good mechanism for this problem.

In practice, there are other very serious issues to handle, like risk-aversion of the agents. While there are practically useful results about it, we cannot expect efficiency or any other mathematically clear result without risk-neutrality. In this paper, we will focus first on the idealized case with risk-neutral agents, but in Appendix E, we will show that these existing risk-sharing techniques are compatible with our offer-based mechanisms.

Problem B. *Our model with trivial capability trees: combinatorial auction.*

Consider the case when the principal has to choose multiple agents and distribute the tasks of the project between them, and completion is the only possible result of each task. This is the combinatorial auction problem with negative-valued goods.

The second-price combinatorial auction means that every agent reports his costs of completing each subset of tasks, the principal chooses the efficient allocation of the tasks calculated from these reports, and each agent gets paid his reported costs plus the second-price compensations. This mechanism implements the socially efficient outcome in dominant strategies. However, this mechanism is extremely vulnerable to collusion if the colluding agents can submit bids which are useless alone but useful together, see Appendix A.3. In contrast, the first-price combinatorial auction is collusion-resistant in the sense that forming a consortium and submitting a joint bid can be beneficial, but bidding separately with a coordinated bid is never better than forming a consortium. Another issue is that the first-price combinatorial auction is individually rational for all players, but the second-price combinatorial auction is not (always) individually rational for the principal.

The first-price combinatorial auction is commonly used in practice, and this is believed to be a reasonably good mechanism, despite the fact that we can prove almost nothing about the level of its efficiency or the expected revenue in the general case. The two simple results about special cases worth mentioning are that the first-price combinatorial auction is efficient (1) if the types are common knowledge, or (2) under perfect competition.

Problem AB. *Our model.*

This problem generalizes both Problems A and B, but it also includes much harder difficulties. As we have already seen, the efficient joint strategy profile means here a high-level of cooperation. Namely, every agent should reveal all of his true abilities and chance events (e.g. faster or slower progress) and they should make the efficient decisions according to the entire current reported state of the project (e.g. always choosing the currently desired effort level). But they can lie and cheat without the risk of being caught. Moreover, in our full model, we will allow all trees to produce contractible consequences which affects the working processes of each other, and we will allow the principal to also have a capability tree. While we clearly cannot avoid the issues we had in Problems A and B, we will essentially resolve the much more serious problem of cooperation failures.

We will show that the second-price mechanism implements efficiency in quasi-dominant equilibrium, although it is vulnerable to collusion. The first-price mechanism is collusion-resistant and works reasonably efficiently. However, as we saw with Problems A and B, it has to be very hard to tell that in what extent it is efficient or provide high utility for the principal. The main message is that this mechanism eliminates the potential losses due to cooperation failures or moral hazard, and it works as well as a first-price combinatorial auction. We will support this message firstly by proving that it works efficiently under special cases: under perfect competition or if the capability trees are common knowledge. Secondly, we will analyse the general case focusing on realistic situations with a reasonably strong competition between the agents.

For practice, we recommend using the first-price offer-based mechanism extended by the practical observations shown in Appendix E.

Problem C. *Our model but with no option of excluding players from participation: dynamic mechanism design, weaker results with very limited applicability.*

There are a number of earlier papers about dynamic mechanism design. The common difference of our model from all of them is that we consider tendering problems, or equivalently, we assume that there is a special player called the principal with the power to keep out any of the other players called agents from the game. In return for this assumption, we got much stronger results.

In more detail, the related literature is the following. Earlier results in dynamic mechanism design started with *online mechanisms* by Friedman and Parkes (2003) [4], and Parkes and Singh (2003) [6], where agents can arrive and depart during the game with hidden utility functions. Cavallo, Parkes, and Singh (2006) [3] proposed a Markovian environment for general allocation problems.

The most closely related results to ours are presented by Athey and Segal (2013) [1] and Bergemann and Välimäki (2010) [2]. Both consider environments in which all players (agents) have the same role, they may have evolving private information and introduce mechanisms which implement the efficient strategy profile in Perfect Bayesian equilibrium. In contrast to our continuous-time model, these models use discrete periods of time in a way that makes the two models importantly different. We assume a complete chronological ordering of the chance events. Namely, we assume that for any two chance events, the report of the earlier one of them cannot depend on the later event. In contrast, they considered the possibility of same-round chance events, namely, some players mutually observe the chance events (stochastic changes of private states) of each other by the time when they report their own chance events. This extra assumption in our model makes it possible to find a mechanism that implements efficiency in a stronger equilibrium and has further important features.

Athey and Segal implement socially efficient decision rules by giving a transfer to each agent that equals the sum of the other agents' flow utilities. This works even if the agents' private signals are correlated. But with independent signals across the agents, they also provide a method of converting any incentive-compatible mechanism into a budget-balanced one. The equilibrium they propose have important weaknesses. It is not always unique, but there can also be inefficient equilibria, moreover, this mechanism is very vulnerable to collusion. Furthermore, the model assumes fixed starting states (initial types). Appendix B describes a lossy translation of our result to their language, and Appendix B.1 describes a detailed comparison.

Bergemann and Välimäki proposed a dynamic pivot mechanism where each agent gets a reward equal to his flow marginal contribution to the total utility. They assume independent signals across agents. Their mechanism has individual rationality in a generalized sense. But their environment does not have private decisions, which are essential in our model. Furthermore, their equilibrium is not guaranteed to be unique, and the mechanism is vulnerable to collusion. Detailed comparison can be found in Appendix B.2.

4 Definitions, the model and the goals

This section formalizes and generalizes the problem described by the example in Section 2. We expect the reader to **understand the example** in detail **before reading this section**. On the other hand, if you do not completely understand something in this section, then you can go on leaning on the understanding of the examples, and come back later to the skipped part.

Section 4.1 is just about clarification of the concepts and the notation we will use. You can even skip this for the first time, and come back to it when some clarification is required.

Section 4.2 defines and generalizes the capability tree of an agent, including the possibilities of interdependencies like precedencies between subtasks of different agents or sharing a common resource. This will make us ready to define the model of the entire project in Section 4.3.

The index of the most important terms at the last page of this paper might be useful.

4.1 General notations and definitions

For any set S and any symbol x , whenever we define x_i for all $i \in S$, then \mathbf{x}_S denotes the vector $(x_i)_{i \in S}$, and \mathbf{x} denotes \mathbf{x}_S for the largest set S for which it is defined. $\mathbf{x}_S \in \mathbf{R}_S$ means $\forall i \in S: x_i \in R_i$. The pair $(\mathbf{x}_S, \mathbf{y}_S)$ is identified with $((x_i, y_i))_{i \in S}$. If $f: A \rightarrow B$ and $\mathbf{x}_S \in A^S$, then $\mathbf{f}(\mathbf{x}_S)$ denotes $(f(x_i))_{i \in S}$. \mathbf{x}_{-j} means the vector of all x_i except x_j , and (y, \mathbf{x}_{-j}) means the vector \mathbf{x} by exchanging x_j to y . The power set of X is denoted by $2^X = \{Y : Y \subseteq X\}$. Game means a potentially dynamic game with an arbitrary ordered set of time points. We will distinguish between information and belief, in order to be able to use belief-independence. In Appendix A.2 we show the reason why this difficulty could not be avoided by simpler techniques.

The following definition clarifies how we will use the terms emphasized by bold text. The notation $(N^*, T_<, A, \alpha, \tilde{\alpha}, \mathcal{I}, \phi, \tilde{\phi}, \tau, \mathcal{B}, \beta, \xi, \sigma)$ will be used in this subsection only.

Definition 1. A **game form** is a tuple $(N^*, T_<, A, \alpha, \mathcal{I}, \phi, \mathcal{B}, \beta, \mathcal{S}, \sigma)$, where

- N^* is the set of strategic and non-strategic **players**,
- $T_<$ is the ordered set of **time points**,
- A is a set of possible actions,
- $\xi_t \in A^{N^* \times \{t' \in T: t' < t\}}$ is the **history** at a time point t , where $\xi_t(i, t')$ is the action made by player i at time point t' ,
- $\alpha: N^* \times \{(t, \xi_t): t \in T, \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}\} \rightarrow 2^A$ defines the **action set**, namely, $\alpha(i, t, \xi_t)$ is the set of feasible actions of player i at time point t after the history ξ_t ,
- \mathcal{I} is the set of possible information (or we could have called them “information identifiers”),
- $\phi: N^* \times \{(t, \xi_t): t \in T, \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}\} \rightarrow \mathcal{I}$ defines the **information** of a player at a time point. This includes the current time and the action set of the agent, namely, there exist functions $\tilde{\alpha}: \mathcal{I} \rightarrow 2^A$ and $\tau: \mathcal{I} \rightarrow T$ such that

$$\forall i \in N^*, \forall t \in T, \forall \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}: \quad \tilde{\alpha}(\phi(i, t, \xi_t)) = \alpha(i, t, \xi_t), \quad \tau(\phi(i, t, \xi_t)) = t.$$

- \mathcal{B} is the set of possible beliefs (or we could have called them “belief identifiers”),
- $\beta: N^* \times \{(t, \xi_t): t \in T, \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}\} \rightarrow \mathcal{B}$ defines the **belief** of a player at a time point, which includes his information, namely, there exists a $\tilde{\phi}: \mathcal{B} \rightarrow \mathcal{I}$ satisfying

$$\forall i \in N^*, \forall t \in T, \forall \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}: \quad \tilde{\phi}(\beta(i, t, \xi_t)) = \phi(i, t, \xi_t),$$

- $\mathcal{S}_i \subset A^{\mathcal{B}}$ is the set of feasible (pure) **strategies**, with a sigma-algebra σ_i on \mathcal{S}_i . A strategy always chooses an action from the feasible action set, namely

$$\forall i \in N^*, \forall s_i \in \mathcal{S}_i, \forall t \in T, \forall \xi_t \in A^{N^* \times \{t' \in T: t' < t\}}: \quad s_i(\beta(i, t, \xi_t)) \in \alpha(i, t, \xi_t).$$

Furthermore, each strategy profile uniquely determines the **execution** ξ , namely,

$$\forall \mathbf{s} = \mathbf{s}_{N^*} \in \mathcal{S}_{N^*}, \exists! \xi: N^* \times T \rightarrow A, \forall i \in N^*, \forall t \in T:$$

$$s_i\left(\beta(i, t, \xi|_{N^* \times \{t' \in T: t' < t\}})\right) = \xi(i, t). \quad (1)$$

The set of mixed strategies of i are the probability distributions of $(\mathcal{S}_i, \sigma_i)$. This σ_{N^*} also induces a sigma-algebra on $A^{N^* \times T}$.

A **game** is a game form with a measurable utility function $u : N \times A^{N^* \times T} \rightarrow \mathbb{R}$. This assigns a real utility $u_i(\xi)$ to each strategic player $i \in N \subseteq N^*$ and each execution ξ .

We note that (1) is not automatic, e.g. this could fail in continuous-time games. This is the reason why we allowed restrictions on the strategy sets (i.e. $\mathcal{S}_i \subset A^{\mathcal{B}}$ and not $\mathcal{S}_i = A^{\mathcal{B}}$). There are different ways to add such restrictions, but there is no universal one. In this paper, any restriction would work as long as this implies (1) and allows some specific simple strategies, e.g. the truthful strategy.

Subgame form means the game form after some period of the game. We define it only for game forms with perfect information, namely, when ϕ is the identity function. Formally, for a game form with perfect information $(N^*, T_{<}, A, \alpha, \mathcal{S})$, an upper-closed set of time points $T' \subseteq T$ (i.e. with $t' \in T'$, $t \in T$, $t' < t \Rightarrow t \in T'$) and a history $\xi_{T \setminus T'} : N^* \times (T \setminus T') \rightarrow A$, subgame is defined as a tuple $(N^*, T'_{<}, A, \alpha|_{N \times T'}, \mathcal{S}')$ where \mathcal{S}' is the set of strategy profiles which are consistent with the history $\xi_{T \setminus T'}$, namely,

$$\mathbf{s} \in \mathcal{S}' \quad \Leftrightarrow \quad \left(\forall i \in N, \forall t \in T \setminus T' : \quad s_i \left(\beta(i, t, \xi|_{N \times \{t' \in T : t' < t\}}) \right) = \xi(i, t) \right).$$

In subgame form structures (e.g. subgames), further but analogous consistency conditions should apply (e.g. for the utility function $\mathbf{u}'(\xi'(\mathbf{s})) = \mathbf{u}(\xi(\mathbf{s}))$). A subgame “preceding” or “after” a time point t (or an event at t) means the subgame with $T' = \{t' \in T : t' \geq t\}$ or $T' = \{t' \in T : t' > t\}$, respectively. A **state** of a game is a synonym of the subgame of the game with perfect information (from a specific time point, or after an event, etc.).

The term of common knowledge refers to the information (not just the beliefs) of the players. The dependence of any function or relation f on the game Γ is denoted by the form f^Γ , but if Γ is the default game, then we may omit this superindex.

We always assume that the rules of the game are known between the players, so formally, all of our games are of complete information. However, we will in fact use some games with incomplete information transformed into an equivalent game of complete information but with nondeterministic actions by nature. Nondeterministic means that we assume no prior distribution.

4.2 Capability tree

The (initial) type of each player including the principal will be a rich structure, describing his entire working capabilities and preferences. We call them capability trees rather than types. Loosely speaking, the capability tree of a player is a dynamic description of

- what decisions he can make during his work, e.g. his options about which technology he uses, how much effort he makes, or how many people he employs;
- what stochastic feedbacks from his working process he expects to observe, e.g. faster progresses or failures;
- how his task is affected by the consequences (externalities) of other tasks, e.g. completion time of a preceding subtasks, reservation of a common resource;
- what consequences the task provides, e.g. the examples above, or the result (main output) of his entire working process;
- the player’s hidden total costs or benefits.

The working process of i can be influenced through contractible consequences of others. But we want to define the working capabilities independently from the other players. This is the reason why we include an abstract player called the world, who chooses the consequences by others.

Definition 2. *Capability tree* is defined as an arbitrary 3-player game form (or 2-player stochastic game form) with a consequence function and a valuation function, which satisfies the following properties.

- The players are: the worker, the world and nature. (They personalize the working efforts of player i , the externalities of others to this capability tree, and the chance events, respectively.)
- The set of time points is I , the same set for all capability trees.
- At each time point in I , the action set of the world is a globally fixed set \mathcal{C} , where $\emptyset \in \mathcal{C}$ represents doing nothing, and if $X, Y \in \mathcal{C}$, then $X \cup Y \in \mathcal{C}$. (The action sets of the worker and nature have no restrictions. These can be arbitrary functions of the earlier moves of all three players, at each time point.)
- Nature makes its moves with a probability distribution determined by the earlier moves of the three players of the capability tree. These are called the **chance events**, at a **chance node** of the capability tree.
- The consequence of the execution of the capability tree at each time point in I is an element of \mathcal{C} determined by the earlier moves of the three players of the capability tree.
- The valuation function $v_i = v(\xi_i)$ assigns a real number (money) to the **execution** ξ_i . (The execution of the capability tree means the entire history of the moves of the three players.)

We say that a capability tree is **omissible** (or idle-allowing) if it has the following properties.

- The worker has a strategy that provides no consequence (provides consequence \emptyset at all points in time), and ends with valuation 0, whatever nature and the world do.
- If any joint strategy of the worker and the world provides no consequence with probability 1, then this ends with expected valuation at most 0.

For example, the capability trees of the players in Example 1 (Section 2) can be defined in the following way.

Example 1	capability tree: game form	consequences	valuation
principal	only the world makes actions	\emptyset in all time points	if the union of the moves of the world was $\{\checkmark_1, \checkmark_2\}$, then 100, otherwise 0
Agents A, B	Figure 3 without results and costs at the leaves; the world does nothing	before the end: \emptyset , at the end: \checkmark_1 or \times_1	the negative of the cost at the leaf
Agents C, D		before the end: \emptyset , at the end: \checkmark_2 or \times_2	

To be more precise, to the capability tree of each agent in Example 1, we should include an option of doing nothing, always providing consequence \emptyset and having valuation 0. Hereby the capability tree of each agent will be omissible. When the principal rejects an agent, this means that she asks the agent to follow this do-nothing strategy, and she applies a huge punishment if he provides any non-empty consequence.

We define the *empty capability tree*, denoted by $[\emptyset]$ to be the capability tree which always provides consequence \emptyset , and the valuation is constantly 0. Having an empty capability tree means having no capability to do anything.

A capability tree of an agent should be interpreted so as it describes the dynamics of what the agent knows about his working process. This is why it is natural and not restricting to say that each agent has perfect information about his own capability tree.

4.3 The model

The primary motivation of the following definition is to describe a general environment where we want to get a particular project completed by some agents. We divide the owner of the project to a planner and a principal. The planner personalizes the behavior of the owner of the project that she can commit to, and the principal personalizes the strategic owner.

We define the Project Management Model, or in short, the **project** denoted by Γ as a game satisfying the following conditions.

There is a player 0 called the **principal** and players $1, 2, \dots, n$ called **agents**. The set of all strategic players is $N = \{0, 1, 2, \dots, n\}$, and the set of agents is $N^+ = \{1, 2, \dots, n\}$. There are two non-strategic players: the **planner** and **nature**. From now on, the word “*player*” will refer to the strategic players, and \mathcal{S} will denote the strategy profile of the players \mathcal{S}_N . Some of the actions of nature are chosen with given probabilities (stochastic moves), and we have no assumption on some other actions (nondeterministic moves). The planner commits to a strategy in advance.

The actions of the players and of nature include the following.

- At a time point t_0 before I , nature arbitrarily assigns a capability tree $\theta_i \in \Theta$ to each player $i \in N$. (Nondeterministic action, no prior.) The agents get omissible capability trees.¹
- The players can send arbitrary contractible time-stamped instant messages to the planner at each real time point after t_0 , and vice versa.
- The capability trees of all players are concurrently executed during I , where each player controls the worker of his capability tree. The actions of the world in each capability tree are defined as the union of the same-time consequences of other capability trees. Each move of nature in a capability tree is made with the given probabilities, independently of everything else up to the current time.²
- At the end (at a time point after I), the planner determines the **payment** (or transfer) vector $\mathbf{p} \in \mathbb{R}^N$ as a function of the contractible events. This must be balanced (including the principal):

$$\sum_{i \in N} p_i = 0. \quad (2)$$

¹Or some of the agents may get non-ommissible but verifiable capability trees, see Section 6.3.

²If a chance event is contractible (e.g. currency exchange rates are of this kind), then we can allow dependence of this chance event with the same-time messages of nature to other players, see Appendix E.3

The information of each player consists of his perfect information in his capability tree and all messages he receives. Namely:

- He has perfect information in his own capability tree. In detail, he knows his capability tree, all strictly earlier moves of the three players in his capability tree, and his chance event at the current time.³
- His information contains all messages he received from the planner strictly previously.

The belief of each player at each time point is assumed to be independent from the concurrent chance events of other players, conditional on the entire past.⁴ This expresses that each player can report the outcome of each of his chance events before any correlated event happens outside of the capability tree. This is our only assumption about beliefs.

By default, the information of the planner is the consequences by all players and the messages she receives until the current time. As a non-default version of the model, the planner may be informed also about the capability tree of the principal. (The planner will not use beliefs. The information and belief of nature is irrelevant as long as it can choose the outcome of the chance events with the desired probabilities.)

The utility \mathbf{u} of the players are defined as follows.

$$\forall i \in \mathbb{N}: \quad u_i = v_i + p_i \quad (3)$$

Only for mathematical accuracy, we need some technical restrictions in order to avoid problems with our continuous-time setting, see in Appendix 4.1.

The subgame after nature assigns the capability tree vector $\boldsymbol{\theta}$ to the players is denoted by $\Gamma(\boldsymbol{\theta})$. Most expressions will be functions of the capability trees as nondeterministic parameters, but when there is no ambiguity, we will drop these parameters.

Interpretation. The owner of the project should agree with the agents about the payment rules, namely, how the payments depend on the achievements (consequences) of the players and on the messages they send to each other during the work. Agreement is optional: on one hand, individual rationality of the agents will mean that they are free to say no and to use their do-nothing strategy. On the other hand, the planner will be able to reject each agent by offering a huge punishment in case if his capability tree provides any consequence. Then all agents who agreed with the owner, as well as the owner concurrently execute their capability trees. At the end, the agents are paid according to the agreements. These payments must depend only on the consequences of the working processes and the communication. Hidden efforts, chance events and other privately observable things of the capability tree cannot directly affect the payment.

Using no prior about the capability trees of the players expresses that we are looking for ex post equilibria with respect to these capability trees. Similarly, we did not specify the beliefs of the players because our solution concept will be robust to them. E.g. the players may be able to make actions that modify the beliefs of others, or they may have chance events which are observable to multiple players. Therefore, the Project Management Model includes environments when different forms of communications like signaling and cheap talk are also possible, or when a mediator is present, etc.

³This is just an unimportant condition for mathematical convenience that we identify the time point of the move of nature with the first time point when the player can react to it. With the formalism introduced in Section 4.1, this means that the set of time points is (a subset of) $\mathbb{R} \times \{0, 1\}$ with the lexicographical ordering, where nature makes actions only at $\{(t, 0) : t \in \mathbb{R}\}$, and the players make actions only at $\{(t, 1) : t \in \mathbb{R}\}$.

⁴Conditional independence between beliefs and *later* chance events is already implied by the independence assumption on the moves of nature.

4.4 Goal

For any $S \subset N$, let $u_S = \sum_{i \in S} u_i$, and we call u_N the **total utility**. Clearly,

$$u_N \stackrel{(2)(3)}{=} \sum_{i \in N} v_i. \quad (4)$$

For any subgame H of Γ , let the (cooperative) value $V(G)$ be the maximum achievable expected total utility in G , namely,

$$V(G) = \sup_{\mathbf{s} \in \mathcal{S}^G} \mathbb{E}(u_N^G(\mathbf{s})). \quad (5)$$

Clearly, for all $\mathbf{s} \in \mathcal{S}$,

$$\mathbb{E}(u_N(\mathbf{s})) \leq V(\Gamma). \quad (6)$$

We want to design a mechanism under which a strategy profile \mathbf{s}^* satisfies the following goals.

- *Efficiency (if all players are risk-neutral):* $\mathbb{E}(u_N(\mathbf{s}^*)) = V(\Gamma)$
- *Incentive compatibility:* \mathbf{s}^* is a “convincing” equilibrium (see Section 8)
- *Collusion-resistance* (see Section 6.2)
- *Individual rationality (or optional participation) for each agent:* at the beginning, if an agent i sends a message to the planner that he does not participate, and then he provides no consequence, then $t_i = 0$.
- *Offers no free lunch:* $\forall i \in N^+$, if $\theta_i = 0$, then $\mathbb{E}(u_i(\mathbf{s}^*)) = 0$,

We will achieve these goals except that either we have to give up collusion-resistance, or we will achieve weaker goals about efficiency.

5 The mechanisms

A **mechanism** is identified with the strategy of the planner. We define two versions of the mechanism, because there is a tradeoff between them. Their relation will be analogous to the relation between the first and second-price auctions (see Section 3).

The true mechanisms called the offer-based mechanisms will be defined in Section 5.3. We define first the revelation mechanism versions of them, because these are easier to handle and to understand, but the offer-based mechanisms will be more useful and more interesting.

5.1 The first-price revelation mechanism

The first-price revelation mechanism is the following. The planner asks the players to report (a claim about) the current state of their capability trees throughout the game. The planner believes these reports, and she always makes and asks the players to make the particular moves in order to maximize the expected total utility. At the end, the payments are determined as follows. The principal pays to each agent the reported cost of the agent, plus whenever a reported chance event of an agent modifies the expected total utility, the principal pays him this (signed) difference.

More formally, the players should use the communication protocol below. At time point t_0 , the principal sends a message to the planner about a capability tree and each agent $i \in N^+$

sends a message to the planner about an omissible capability tree. We call these messages of the agents their **proposals** $\hat{\theta}_i$ (which can be different from his actual capability tree θ_i). Then, the planner governs a strategy profile $\hat{\mathbf{s}}$ of the **reported project** $\Gamma(\hat{\boldsymbol{\theta}})$ which maximizes $\mathbf{E}(u_{\mathbf{N}}^{\Gamma(\hat{\boldsymbol{\theta}})})$. In detail, she chooses an

$$\hat{\mathbf{s}} \in \arg \max_{\mathbf{s} \in \mathcal{S}^{\Gamma(\hat{\boldsymbol{\theta}})}} \mathbf{E}(u_{\mathbf{N}}^{\Gamma(\hat{\boldsymbol{\theta}})}(\mathbf{s})), \quad (7)$$

- and whenever each player $i \in \mathbf{N}$ reaches a chance node in $\Gamma(\hat{\boldsymbol{\theta}})$, the planner requires him to report her an outcome of the chance event, and she considers the reported outcome to be the outcome of the chance event in $\Gamma(\hat{\boldsymbol{\theta}})$. For any two chance events, we consider them to occur at distinct points in time. In a case of concurrence, we just define an arbitrary ordering of the chance events.
- Whenever an agent $i \in \mathbf{N}$ is about to make a move in $\Gamma(\hat{\boldsymbol{\theta}})$, the planner sends him the move corresponding to \hat{s}_i , and she considers this to be his move in $\Gamma(\hat{\boldsymbol{\theta}})$.
- At the end, the planner gets to a reported execution $\hat{\xi}_i$ of each $i \in \mathbf{N}$. If the actual consequences provided by an agent $i \in \mathbf{N}^+$ were different from the consequences indicated by his reported execution $\hat{\xi}_i$, then $p_i = -\infty$ (or a very big negative value). Otherwise, denote the set of the chance events during the reported execution of $\hat{\theta}_i$ in $\Gamma(\hat{\boldsymbol{\theta}})$ by X_i , and for each $\chi \in X_i$, denote the states of the project preceding χ by T^χ and succeeding χ by T_+^χ . Let the change of the expected total utility by χ be denoted by

$$\delta(\chi) = V(T_+^\chi) - V(T^\chi). \quad (8)$$

Then the planner chooses the following payment to each agent i .

$$p_i = -\hat{v}_i(\hat{\xi}_i) + \sum_{\chi \in X_i} \delta(\chi), \quad (9)$$

and therefore, the payment of the principal is

$$p_0 = - \sum_{i \in \mathbf{N}^+} p_i. \quad (10)$$

The punishment $p_i = -\infty$ expresses that the report of an agent about his capability tree and the execution of his capability tree may be completely different from the reality, but it must match with the observable reality. In other words, even if an agent lies, he must use a consistent explanation of how he has provided his consequences. From now on, we assume that each agent i uses a strategy that guarantees that this punishment will not be applied. Or we can use the technique in Appendix E.1.

We note that the “value of a state of the project” in Examples 1 and 2 corresponds to $V(T)$ here.

5.2 The second-price revelation mechanism

We define the **joint value of capability trees** $\boldsymbol{\theta}'$ by $V(\boldsymbol{\theta}') = V(\Gamma(\boldsymbol{\theta}'))$. The **value of a proposal** $\hat{\theta}_i$ is the reported marginal contribution

$$v_i^+ = v_i^+(\hat{\boldsymbol{\theta}}) = V(\hat{\boldsymbol{\theta}}) - V(\hat{\boldsymbol{\theta}}_{-i}). \quad (11)$$

The second-price revelation mechanism is the same as the first-price revelation mechanism except that the principal pays v_i^+ more to each agent $i \in N^+$, namely,

$$p_i^{2nd} = p_i + v_i^+(\hat{\theta}) \stackrel{(9)}{=} v_i^+(\hat{\theta}) - \hat{v}_i(\hat{\xi}_i) + \sum_{\chi \in X_i} \delta(\chi).$$

(To avoid ambiguity of notations, we denote the payment to $i \in N$ under the second-price revelation mechanism by p_i^{2nd} .)

If we use the assumption that the planner knows the capability tree of the principal, then the principal does not have to report it, but we define $\hat{\theta}_0 = \theta_0$.

5.3 The offer-based mechanisms

First, we apply the following modification of the mechanisms. Instead of using communication between the planner and the players, principal can send and receive all these messages, but the planner sees them, and she always considers the decisions of the players in $M(\hat{\theta})$ to be the suggestion sent by the principal. Now, the truthful strategy of the principal includes that she always makes suggestions that maximizes the expected total utility. All proofs are valid in this case, as well. (Except that the proofs of the weaker version of the results in Section 7 require some modifications.)

We define the first-price offer-based mechanism in the following way. We define the contract as a function that determines the payment between two players depending on the consequences provided by all players and the communication between the two players. Proposal means a contract offered by an agent to the principal. At the beginning, each agent sends a proposal to the principal. Then the principal accepts some of them, and the other agents leave the game with utility 0.⁵ The only role of the planner is that she observes the communication including the proposals, and she determines the payments at the end, according to the accepted contracts.

Now, we define the “conspiracy-fearing” strategy of the principal τ_0 . After the principal receives the proposals, she uses the strategy out of all her possible strategies by which *her minimum possible expected utility is the largest*, in the following sense. After the proposals are submitted, we define the following two-player game. One player is the principal with the same capability tree (and with true chance events). The other player is called the devil who controls the moves of all agents, and we replace their capability trees with universal capability trees, namely, the devil has the capability to provide arbitrary consequences assigned to each agent and to send messages in the name of each agent, and he observes the capability tree of the principal as well as its execution (excluding her same-time chance events). The payments at the end are determined according to the contracts. The aim of the devil is to minimize the expected utility of the principal, while the principal aims to maximize it. The principal has a maximin strategy in this two-player zero-sum game. The conspiracy-fearing strategy of the principal in our original game is defined as she makes the very same moves as what she would do by the maximin strategy in this game against the devil. We call the value of this two-player game the **maximin utility** of the principal. Correspondingly, we define the joint value of a set of proposals as her maximin utility if she accepts this set of proposals.

Let us see first how it works with a simple single-agent task. We want to hire somebody for a task, for simplicity, assume that the result can only be of high or medium quality. The difference is equivalent to \$100 for us.

⁵Rejecting an agent means that he is asked to do nothing, and if he provides any consequence, then he should pay a large punishment (as mentioned in Section 4.3 – interpretation).

Assume that an agent A would ask for \$200 and would provide high quality with probability 50% and medium quality with probability 50%. If he really asks for \$200, then with our untrustful (“one-agent conspiracy-fearing”) strategy, we do not believe the reported probabilities but we expect a medium quality result from A . This means that we would still prefer an agent B who offers to make the task for \$199 in medium quality. Therefore, A should make the following offer instead. In case if he completes the task in medium quality, then he asks for \$150, but if in high quality, then he asks for \$250. In this case, the two results are equally good for us, and therefore, we do not need to care about the probabilities, but this offer beats another offer of completing the task even for \$151 in medium quality. So A gets \$200 payment in expectation and this offer is evaluated to be as good as if we believed his probabilities.

The truthful offers of some capability trees are described in Section 2.1. A fair offer is the very same but with payments increased by a constant.

Notice that every proposal in the first-price revelation mechanism can be translated to an equivalent proposal in the offer-based mechanism. In fact, the offer-based mechanism provides a richer set of possible proposals, which can be useful when we want to use the mechanism under more general environments, e.g. with risk-aversion, limited responsibility, etc, see Appendix E.

Analogously, we could define the second-price offer-based mechanism as the same mechanism but the principal pays an extra second-price compensation v_i^+ money to each agent i beyond the payment according to their contract. v_i^+ is defined as the marginal increment of the maximin utility of the principal by the proposal of the agent. Here, we assume again that the planner can observe the capability tree of the principal (but not the execution).

All results remain true with almost the same proof under this mechanism, but we will elaborate more on it in Appendix C.

5.4 Definitions and notations

By default, we will use the revelation mechanisms just because of their simpler syntax, but the results will be valid also for the offer-based mechanisms, or we will address the differences. Let Γ/m denote the game Γ given the mechanism (strategy of the planner) m . Let m_1 and m_2 denote the first and the second-price (revelation) mechanisms, respectively. Let $M = \Gamma/m_1$ be the default game and all notions defined in Γ are used correspondingly in M , and clearly,⁶ $V(\Gamma) = V(M)$. The game forms of the games $M = \Gamma/m_1$ and Γ/m_2 are the same; the only difference between them is the utility function. Therefore, we can define the second-price utilities as functions in M , and we denote them by u^{2nd} . This means that

$$\forall i \in N^+ : \quad u_i^{2nd} = u_i + v_i^+(\hat{\theta}), \quad (12)$$

$$u_0^{2nd} = u_0 - \sum_{i \in N^+} v_i^+(\hat{\theta}). \quad (13)$$

We define the **truthful strategy** τ_i of a player $i \in N$ as follows. He reports his true capabilities, namely, $\hat{\theta}_i = \theta_i$, and then he always makes the move which the planner asks him, and at each chance node, he reports the true outcome of the chance event.

6 Results

6.1 Second-price mechanism

Despite the fact that no dominant strategy equilibrium exists in such a general model (see Appendix A.1), we implement the efficient strategy profile in a slightly weaker equilibrium

⁶because of the actions of the planner does not decrease the maximum possible total utility of the players

called quasi-dominant equilibrium. This definition is very difficult and unusual, we do not expect the reader to see the point at the first glance, but we will spend the entire Section 8 on this equilibrium concept. We will show there that it is a stronger concept than Perfect Bayesian equilibrium in a reasonable sense, and also stronger than roughly every such refinement of Nash-equilibrium which exists in all games.

Definition 3. *There is a set of players N playing a stochastic dynamic game with imperfect information. \mathcal{S}_i denotes the strategy set and $u_i : \mathcal{S} \times \Theta^N \rightarrow \mathbb{R}$ denotes the utility function of player $i \in N$. At the very beginning, nature assigns a(n initial) type $\theta_i \in \Theta$ for each player $i \in N$, as a nondeterministic assignment with no assumed prior probabilities. Each player gets to know only his own type. Then at the same time point, the players do their first actions simultaneously, denoted by $a_i = \alpha(\theta_i, s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \Theta \times \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $\alpha^*(\theta_i) = \alpha(\theta_i, s_i^*)$, satisfying the following.⁷*

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i(\boldsymbol{\theta}, (s_i^*, \mathbf{s}_{-i}))\right) \geq f_i(\theta_i, \boldsymbol{\alpha}(\boldsymbol{\theta}_{-i}, \mathbf{s}_{-i}))$$

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall a_i \in \mathcal{A}_i:$$

$$\sum_{j \in N} \mathbb{E}(u_j(\boldsymbol{\theta}, \mathbf{s})) \leq f_i(\theta_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i})) + \sum_{j \in N \setminus \{i\}} f_j(\theta_j, (a_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i-j})))$$

Then \mathbf{s}^* is a **quasi-dominant equilibrium**.

We do not expect this equilibrium concept to be widely used, because it seems extremely rare that dominant strategy equilibrium is impossible but quasi-dominant strategy implementation exists. Nevertheless, this is the case in our model:

Theorem 1. *If the planner knows the capability tree of the principal, then the truthful strategy profile $\boldsymbol{\tau}$ is a quasi-dominant equilibrium under the second-price mechanism.*

The proof, presented in Section 9.1, is very abstract, but there is a much more intuitive proof of the much weaker concept of Nash-equilibrium in Section 7.

The second-price mechanism is individually rational for the agents, because if an agent $i \in N^+$ reports the empty capability tree and chooses the do-nothing option, then his payment is $p_i = 0$. Also, the mechanism does not offer free lunch, because if agent i has an empty capability tree and he is truthful, then again, he gets payment $p_i = 0$. Both arguments are valid under the first-price mechanism, as well.

Furthermore, we conjecture that the principal does not benefit from the ability of verifiably showing her capability tree (the planner knows her capability tree), see Conjecture 23.

To sum up the claimed results, the second-price mechanism satisfies all of our goals except individual rationality for the principal, and collusion resistance. Indeed, this mechanism does not achieve these two goals. Both problems arises from the second-price compensations, as shown in the example in Appendix A.3. We note that this problem is present even in a second-price combinatorial auction, which is clearly a special case of our Project Management Model with the second-price mechanism.

⁷Remember that with our notation, $\boldsymbol{\alpha}(\boldsymbol{\theta}_{-i}, \mathbf{s}_{-i}) = (\alpha(\theta_j, s_j))_{j \in N \setminus \{i\}}$ and $\boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i-j}) = (\alpha^*(\theta_k))_{k \in N \setminus \{i,j\}}$ and $\boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i}) = (\alpha^*(\theta_j))_{j \in N \setminus i}$.

6.2 First-price mechanism

In short, the first-price mechanism offers individual rationality also for the principal, it does not assume that the planner knows the capability tree of the principal, and what is more important, it is collusion resistant. In exchange, this implements only an approximately efficient outcome.

In the Project Management Model, selfish private decisions and false reports can cause much more serious losses than inefficient allocation of the tasks to the agents. Therefore, it is a reasonable and ambitious goal to present a mechanism for the Project Management Model with the same properties as those of the first-price auctions in Problems A and B in Section 3.

6.2.1 Individual rationality for the principal

The following theorem proves individual rationality for the principal, namely, she gets at least as much utility as what she could get alone in the game $M(\theta_0)$.

Theorem 2.

$$\mathbb{E}(u_0(\tau)) \geq \mathbb{E}(u_0^{M(\theta_0)}(\tau_0))$$

Proof. Proposition 9 says that the expected utility of the truthful principal is always equal to the value of the reported project, namely, $\mathbb{E}(u_0(\tau_0)) = V(M(\hat{\theta}))$. But $M(\hat{\theta})$ contains the possibility that all agents do nothing. This essentially implies the theorem. Formally, let \emptyset denote the strategy of reporting the empty capability tree $[\emptyset]$ and then using the do-nothing strategy.

$$\mathbb{E}(u_0(\tau)) \stackrel{(18)}{=} V(M(\hat{\theta})) \stackrel{(5)}{\geq} \mathbb{E}(u_N^{M(\hat{\theta})}(\tau_0, \emptyset^{N^+})) = \mathbb{E}(u_N^{M(\hat{\theta}_0)}(\tau_0)) = \mathbb{E}(u_0^{M(\hat{\theta}_0)}(\tau_0)) = \mathbb{E}(u_0^{M(\theta_0)}(\tau_0)). \quad \square$$

6.2.2 Collusion resistance

First, we give an informal proof that colluding is not better than forming a consortium.

Assume that an agent has multiple capability trees (and plays as the worker in each of them at the same time). This can be equivalently described as if he had the “combined” capability tree (like the combined trees in Examples 1 and 2), which is the capability tree with the product of the action sets, the union of the consequences and the sum of the corresponding utilities.

Now, assume that some players are colluding. Let us replace these agents with a new agent with the combined capability tree. Any joint strategy of the colluding players corresponds to a strategy of the new player, and this player gets the same expected utility as the expected total utility of the coalition had been. This provides a reduction of the problem with colluding players to the original problem.

More formally, we define the combination (or the product) of some capability trees $\theta_x = \prod_{i \in X} \theta_i$ as the following capability tree. The three players simultaneously play the games θ_i , making the moves in chronological order. Nature uses the strategies specified in the capability trees θ_i , independently. At each time point, if the action of the world in the combined capability tree θ_x is $c \in \mathcal{C}$ and the consequence provided by each capability tree θ_j is c_j , then the action of the world in θ_i is defined as $c \cup \bigcup_{j \in X \setminus \{i\}} c_j$. The worker of the combined capability tree θ_x plays as the worker in each θ_i . The consequence of θ_x is defined as $\bigcup_{i \in X} c_i$. The valuation of an execution of θ_x is the sum of the valuations of the individual capability trees, namely for an execution $\xi_X = (\xi_i)_{i \in X}$, we define $v(\xi_X) = \sum_{i \in X} v(\xi_i)$.

Lemma 3. For a subset of agents $X \subset \mathbb{N}^+$ and strategies $s_i \in \mathcal{S}_i(\theta_i)$, there exists a strategy $s_x \in \mathcal{S}_x\left(\prod_{i \in X} \theta_i\right)$ satisfying ⁸

$$\sum_{i \in X} \mathbb{E}(u_i^{M(\theta)}(\mathbf{s})) = \mathbb{E}\left(u_x^{M\left(\left(\prod_{i \in X} \theta_i, \boldsymbol{\theta}_{\mathbb{N} \setminus X}\right)\right)}(s_x, \mathbf{s}_{\mathbb{N} \setminus X})\right), \quad (14)$$

where x denotes the agent in $M\left(\left(\prod_{i \in X} \theta_i, \boldsymbol{\theta}_{\mathbb{N} \setminus X}\right)\right)$ with the combined capability tree $\theta_x = \prod_{i \in X} \theta_i$.

Proof. We define the product strategy $\prod_{i \in X} s_i \in \mathcal{S}_x\left(\prod_{i \in X} \theta_i\right)$ of agent x as follows. He sends the proposal $\prod_{i \in X} \hat{\theta}_i(s_i)$, and simulates the case as if he would follow the strategy s_i in each capability tree θ_i , namely, he makes the corresponding actions in the capability trees, and he always reports the product of the states what he would report by each s_i .

Hereby the executions of $M\left(\left(\prod_{i \in X} \theta_i, \boldsymbol{\theta}_{\mathbb{N} \setminus X}\right)\right)$ correspond to the executions of $M(\theta)$ by a natural transformation which preserves probabilities. For each corresponding pair of executions of the project ξ and ξ' ,

$$\sum_{i \in X} u_i^{M(\theta)}(\mathbf{s}, \xi) = u_x^{M\left(\left(\prod_{i \in X} \theta_i, \boldsymbol{\theta}_{\mathbb{N} \setminus X}\right)\right)}\left(\left(\prod_{i \in X} s_i, \mathbf{s}_{\mathbb{N} \setminus X}\right), \xi'\right).$$

Taking expectation for a random execution, we get (14). □

6.2.3 Level of efficiency and revenue maximization

We want to prove the same results that are true for the first-price auction mechanisms. In particular, we prove efficiency in two extreme cases: (1) under perfect competition and (2) with publicly known capability trees (but still with private decisions and hidden chance events). Then we show a justification that in a general case, the players are approximately incentivized to follow essentially truthful strategies, and we prove an inequality which also suggests that it is very exceptional having incentives to an essential deviation from it.

We say that there is **perfect competition** if it is common knowledge between the agents that $\forall i \in \mathbb{N}^+ : V(\boldsymbol{\theta}_{-i}) = V(\boldsymbol{\theta})$.

Theorem 4. *The truthful strategy profile $\boldsymbol{\tau}$ under the first-price mechanism is a quasi-dominant equilibrium under perfect competition.*

The proof is presented in Section 9.2.

Under imperfect competition, analogously to the case of first-price auctions, we expect that the agents ask higher payments – report lower valuations – than their real costs. In this case, we can no longer use a belief-free equilibrium, but we need to use a Bayesian belief system, see in Section 10.

For a capability tree θ_i and a number x (which depends on the belief of the agent when he submits his proposal), let $\theta_i - x$ be the same capability tree but with valuation decreased by a constant x except for the do-nothing execution. In addition, for a strategy s_i using proposal $\hat{\theta}_i$, let $s_i - x$ mean the strategy by which i always makes the same actions as he would do with s_i ,

⁸ $\mathcal{S}_i(\theta)$ is defined as the set of strategies of i restricted to the event that his capability tree is θ . We are using the very technical assumption that $\mathcal{S}_x\left(\prod_{i \in X} \theta_i\right)$ includes all strategies which we will mention.

except that he sends the proposal $\hat{\theta}_i - x$ instead. We call a strategy $\tau_i - x$ the **fair strategy** and a proposal $\theta_i - x$ a **fair proposal** with **profit** x .

Consider the extreme case when the players can observe the capability trees of each other. We emphasize that this assumption does not mean that the players can observe anything about the executions of the capability trees of others.

Theorem 5. *If the capability trees of the players are common knowledge between the agents, then there exists a weak quasi-dominant equilibrium where the agents use fair strategies, the principal uses truthful strategy, they maximize the expected total utility, and the principal gets at least as much utility as under the second-price mechanism (with the quasi-dominant strategy profile).*

The precise definitions, the theorems and proofs are presented in Sections 8.1 and 11.

In Section 10, we will elaborate more about the general case, using a Bayesian belief system, including an inequality suggesting that the agents have very weak incentives in deviating from their best fair strategies. In Appendix F, we introduce an environment for analyzing different versions of general auction mechanisms including the Project Management Model. For example, this could help analyzing the version with “the core compensation closest to the second-price compensation”. This version might be a good tradeoff between efficiency and collusion-resistance, but we will not elaborate on it.

6.3 Non-omissible but verifiable capability trees

Assume that some agents have non-omissible capability trees but these are observed by the planner. We emphasize that the executions of their capability trees are still hidden. All of our results extend to this case with essentially the same proof (except individual rationality for these agents, which is not a reasonable goal for them). We only need the following modifications. Such an agent i does not report his capability tree, but we define $\hat{\theta}_i = \theta_i$. We define $v_i^+ = 0$, namely, he does not get any second-price compensation.

6.4 Overview of the results

We introduced a general model for multi-agent projects when each agent has a separate hidden working process with private actions and hidden chance events, but they have influence on each other through publicly observable actions and events. In other words, this is a dynamic stochastic multi-agent model with private values. Despite the fact that the project requires a high-level of cooperation and the agents can cheat in various ways without the risk of being caught, we have designed a mechanism which incentivizes the players including the principal in truthful behavior, in a very strong sense. We introduced two very similar mechanisms, and there is a tradeoff between them: the first-price mechanism is better in almost all aspects which are important in practice, including collusion-resistance, but this is only approximately efficient. This is more a theoretical than a practical deficiency. The second-price mechanism implements social efficiency in a quasi-dominant equilibrium, but even two colluding players can make it very bad in practice. More formally, the results are the following.

Under the second-price mechanism, the truthful strategy profile is a quasi-dominant equilibrium, and this way the expected total utility is the largest possible. We emphasize that dominant strategy equilibrium does not exist under such a general model, but surprisingly we could implement the efficient strategy profile with just a slightly weaker equilibrium concept.

Under the first-price mechanism, with some assumption (or approximation), there is an equilibrium consisting of the same strategies but demanding money increased by a constant (their expected utilities in the case of acceptance) in the proposals.

Both mechanisms are individually rational for the agents, and useless agents get utility 0.

The *disadvantage* of the first-price mechanism over the second-price mechanism is that it requires competition on each task, and this is completely efficient only under perfect competition or if the capabilities of all agents are commonly known.

The *advantages* of the first-price mechanism are the following.

- Collusion does not decrease efficiency as long as it does not decrease the competition, while under the second-price mechanism, if two agents can send proposals that are useless without each other, then they may get as much utility as they want.
- The capabilities of the principal does not need to be known by the planner.
- The mechanism is individually rational also for the principal.

The first-price mechanism can be interpreted as follows. The principal asks the agents to send her offers for contract, and then the principal chooses which contracts to sign and how to communicate with the agents so as to maximize her minimum possible utility. Surprisingly, this incentivizes the agents to follow the efficient strategy profile.

We will show numerous important details how to use the mechanism in practice, see Appendices D and E.

7 Proofs of the basic goals with Nash-equilibrium

Here, we prove that the second-price mechanism achieves our goals (Section 4.4) with Nash-equilibrium. This is not an adequate equilibrium concept for our dynamic game, but the argument describes the essence of the proof of the stronger equilibrium concepts, and it is much easier to understand. We will improve the results later by showing that this argument can prove much more than a Nash-equilibrium.

In view of Examples 1 and 2, the next five lemmas and Proposition 11 are easy to prove. In particular, Lemmas 6 and 7 show that the values of the states of the project can be calculated in the same way as in Examples 1 and 2. The main idea of this section is in Theorem 12.

Lemma 6. *Denote the state of the project preceding a move of a player by T , and all possible states of the project succeeding the different moves by T_1, T_2, \dots, T_m . Then*

$$V(T) = \max_k V(T_k). \quad (15)$$

Proof. In the subgame T , the strategy profile of the players consists of a move $k \in \{1, 2, \dots, m\}$ and a strategy profile in T_k . Therefore,

$$V(T) \stackrel{(5)}{=} \max_{\mathbf{s} \in \mathcal{S}^T} \mathbf{E}(u_N^T(\mathbf{s})) = \max_k \max_{\mathbf{s} \in \mathcal{S}^{T_k}} \mathbf{E}(u_N^{T_k}(\mathbf{s})) \stackrel{(5)}{=} \max_k V(T_k). \quad \square$$

Lemma 7. *Denote the state of the project at a chance node of a player by T , and all possible states of the project succeeding the chance event by T_1, T_2, \dots, T_m . Let w_1, w_2, \dots, w_m denote the corresponding probabilities. Then $V(T) = \sum w_k \cdot V(T_k)$, or equivalently,*

$$\sum_{k=1}^m w_k (V(T_k) - V(T)) = 0. \quad (16)$$

Proof. There is a natural correspondence between \mathcal{S}^T and $\times_{k=1}^m \mathcal{S}^{T_k}$, namely, each strategy profile in T consists of what the players would do after the different outcomes of the chance event. Thus,

$$V(T) \stackrel{(5)}{=} \max_{\mathbf{s} \in \mathcal{S}^T} \mathbb{E}(u_N^T(\mathbf{s})) = \max_{\mathbf{s} \in \mathcal{S}^T} \sum_{k=1}^m w_k \cdot \mathbb{E}(u_N^{T_k}(\mathbf{s})) = \sum_{k=1}^m w_k \cdot \max_{\mathbf{s} \in \mathcal{S}^{T_k}} \mathbb{E}(u_N^{T_k}(\mathbf{s})) \stackrel{(5)}{=} \sum_{k=1}^m w_k \cdot V(T_k). \quad \square$$

Let the stochastic variable \tilde{X}_i denote the set of chance events of i , given the capability trees and the strategies.

Lemma 8. *Assume that a player follows truthful strategy except maybe misreporting his valuation function. Then,*

$$\mathbb{E}_{\tilde{X}_i} \left(\sum_{\chi \in X_i} \delta(\chi) \right) = 0. \quad (17)$$

Proof. Using the notions in Lemma 7, $\mathbb{E}(\delta(\chi)) \stackrel{(8)}{=} \sum w_i (V(T_i) - V(T)) \stackrel{(16)}{=} 0$. Consequently, $\sum \delta(\chi)$, summing on all past chance events χ of $i \in N$, is a martingale, and therefore, the expected values of the sums until the end (left hand side) and until the beginning (right hand side) are the same. \square

The following lemma shows that the expected utility of the principal depends only on the proposals, no matter how the agents behave afterwards. (If the principal has no chance event, then no expectation is required.) Recall that u_i is the utility with the first-price payments, not with the second-price payments.

Proposition 9. *If the principal is truthful, ie. $s_0 = \tau_0$, then*

$$\mathbb{E}(u_0) = V(\hat{\theta}) \quad (18)$$

Proof. Denote the current reported state of the project by T , the reached reported final state by F , denote the sequence of the reported chance events of $i \in N$ until T by $X_i(T)$, and let $X(T) = \bigcup_{i \in N} X_i(T)$. In particular, $X(M(\hat{\theta})) = \emptyset$. Now $V(T) - \sum_{\chi \in X(T)} \delta(\chi)$ is invariant during the game because of (8) and (15). If we apply this with $T = M(\hat{\theta})$ and $T = F$, we get

$$V(\hat{\theta}) = V(M(\hat{\theta})) = V(M(\hat{\theta})) - \sum_{\chi \in X(M(\hat{\theta}))} \delta(\chi) = V(F) - \sum_{\chi \in X(F)} \delta(\chi). \quad (19)$$

Therefore, using that $\hat{v}_0(\tau_0) = v_0$ and $\hat{\xi}_0(\tau_0) = \xi_0$,

$$\begin{aligned} V(\hat{\theta}) &\stackrel{(19)}{=} \mathbb{E} \left(V(F) - \sum_{\chi \in X(F)} \delta(\chi) \right) \stackrel{(5)(4)}{=} \mathbb{E} \left(\sum_{i \in N} \hat{v}_i(\hat{\xi}_i) \right) \stackrel{(17)}{=} \mathbb{E}(\hat{v}_0(\hat{\xi}_0)) + \sum_{i \in N^+} \mathbb{E} \left(\hat{v}_i(\hat{\xi}_i) - \sum_{\chi \in X_i} \delta(\chi) \right) \\ &\stackrel{(9)}{=} \mathbb{E} \left(v_0(\xi_0) - \sum_{i \in N^+} p_i \right) \stackrel{(3)(10)}{=} \mathbb{E}(u_0). \quad \square \end{aligned}$$

Lemma 10. *The expected utility of each agent who uses truthful strategy is 0, no matter how the other players behave. Formally,*

$$\mathbb{E}(u_i(\tau_i)) = 0 \quad (20)$$

Proof.

$$\mathbb{E}(u_i) \stackrel{(3)}{=} \mathbb{E} \left(v_i(\xi_i) + p_i \right) \stackrel{(9)}{=} \mathbb{E} \left(v_i(\xi_i) - \hat{v}_i(\hat{\xi}_i) + \sum_{\chi \in X_i} \delta(\chi) \right) \stackrel{(17)}{=} \mathbb{E} \left(v_i(\xi_i) - \hat{v}_i(\hat{\xi}_i) \right) = 0. \quad \square$$

The following proposition shows the efficiency of the truthful strategy profile.

Proposition 11.

$$\mathbb{E}(u_N(\boldsymbol{\tau})) = V(M). \quad (21)$$

Proof. If $\mathbf{s} = \boldsymbol{\tau}$, then $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$, therefore,

$$\mathbb{E}(u_N) = \mathbb{E}(u_0) + \sum_{i \in N^+} \mathbb{E}(u_i) \stackrel{(18)(20)}{=} V(\hat{\boldsymbol{\theta}}) = V(\boldsymbol{\theta}) = V(M). \quad \square$$

Theorem 12. *If the planner knows the capability tree of the principal, then the truthful strategy profile $\boldsymbol{\tau}$ is a Nash-equilibrium under the second-price mechanism.*

Proof. Notice that

$$\mathbb{E}(u_0) - v_i^+ \stackrel{(18)(11)}{=} V(\hat{\boldsymbol{\theta}}) - (V(\hat{\boldsymbol{\theta}}) - V(\hat{\boldsymbol{\theta}}_{-i})) = V(\hat{\boldsymbol{\theta}}_{-i}). \quad (22)$$

We show that the truthful strategy profile $\boldsymbol{\tau}$ is a Nash-equilibrium, because $\forall i \in N^+$:

$$\begin{aligned} & \mathbb{E}(u_i^{2nd})(\boldsymbol{\theta}, (s_i, \boldsymbol{\tau}_{-i})) \stackrel{(12)}{=} (\mathbb{E}(u_i) + v_i^+)(\boldsymbol{\theta}, (s_i, \boldsymbol{\tau}_{-i})) \\ &= \left(\mathbb{E}(u_N) - \sum_{j \in N^+ \setminus \{i\}} \mathbb{E}(u_j) - \mathbb{E}(u_0) + v_i^+ \right)(\boldsymbol{\theta}, (s_i, \boldsymbol{\tau}_{-i})) \\ & \stackrel{(6)(20)(22)}{\leq} V(\boldsymbol{\theta}) - \sum_{j \in N^+ \setminus \{i\}} 0 - V(\hat{\boldsymbol{\theta}}_{-i})(\boldsymbol{\theta}, (s_i, \boldsymbol{\tau}_{-i})) = V(\boldsymbol{\theta}) - V(\boldsymbol{\theta}_{-i}), \end{aligned}$$

with equation if $s_i = \tau_i$. Therefore, for all $\boldsymbol{\theta} \in \Theta^N$,

$$\mathbb{E}(u_i^{2nd})(\boldsymbol{\theta}, (s_i, \boldsymbol{\tau}_{-i})) \leq V(\boldsymbol{\theta}) - V(\boldsymbol{\theta}_{-i}) = \mathbb{E}(u_i^{2nd})(\boldsymbol{\theta}, \boldsymbol{\tau}).$$

Consider now the principal.

$$\begin{aligned} & \mathbb{E}(u_0^{2nd})(\boldsymbol{\theta}, (s_0, \boldsymbol{\tau}_{-0})) \stackrel{(13)}{=} \left(\mathbb{E}(u_0) - \sum_{i \in N^+} v_i^+ \right)(\boldsymbol{\theta}, (s_0, \boldsymbol{\tau}_{-0})) \\ &= \mathbb{E}(u_0(s_0, \boldsymbol{\tau}_{-0})) - \sum_{i \in N^+} v_i^+(\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}) \stackrel{(18)}{\leq} V(M) - \sum_{i \in N^+} v_i^+(\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}), \end{aligned}$$

and equation holds if $s_0 = \tau_0$. \square

The following lemma implies that the second-price mechanism offers no free lunch.

Lemma 13. *If $\theta_i = 0$ and $\mathbf{s} = \boldsymbol{\tau}$, then $u_i = u_i^{2nd} = 0$.*

Proof. If $\theta_i = 0$, then $\hat{\theta}_i(\boldsymbol{\tau}) = 0$, therefore, $u_i = 0$. Furthermore, $V(\boldsymbol{\theta}) = V(\boldsymbol{\theta}_{-i})$, therefore,

$$u_i^{2nd} \stackrel{(12)}{=} u_i + v_i^+(\hat{\boldsymbol{\theta}}) \stackrel{(11)}{=} 0 + V(\boldsymbol{\theta}) - V(\boldsymbol{\theta}_{-i}) = 0. \quad \square$$

Summarizing the results, the second-price mechanism is individually rational, offers no free lunch, and the truthful strategy profile is a Nash-equilibrium and maximizes the expected total utility.

8 The quasi-dominant equilibrium

Theorem 12 was about Nash-implementation, but we show that the proof implies a strong kind of incentive compatibility, which works even in dynamic environments. In this section, we define the quasi-dominant equilibrium, which is designed to catch the incentive compatibility the proof provides. Then, in Section 9.1, we will see that the proof indeed works with this equilibrium.

The strength of an equilibrium concept is an interpretational question. However, we try to provide a justification as close to a mathematical proof as possible. Starting with a very special case, we arrive at the quasi-dominant equilibrium and its justification in several steps.

In Section 8.2, we will show that quasi-dominant equilibrium is stronger than Perfect Bayesian equilibrium, in a reasonable sense. These justifications also give a good intuition for that proof.

At all steps, \mathcal{S}_i denotes the strategy set and u_i denotes the utility function of player i .

Step 1. Players A and B play an arbitrary deterministic *dynamic* game with *imperfect information*. Suppose that there is a strategy profile $\mathbf{s}^* \in \mathcal{S}$ satisfying the following.

$$\forall s_B \in \mathcal{S}_B: \quad u_A(s_A^*, s_B) \geq 1 \quad (23)$$

$$\forall s_A \in \mathcal{S}_A: \quad u_B(s_A, s_B^*) \geq 1 \quad (24)$$

$$\forall \mathbf{s} \in \mathcal{S}: \quad u_A(\mathbf{s}) + u_B(\mathbf{s}) \leq 2 \quad (25)$$

Then \mathbf{s}^* is an equilibrium.

Justification. In an arbitrary two-player game, if either player can guarantee himself utility 1, and even with collusion, they cannot get more utility in total, then both players will use the strategy guaranteeing utility 1.

More formally, (23) shows that A can get utility at least 1, therefore, if she is selfish and rational, then she will get expected utility at least 1. Comparing this with (25), this shows that B has no hope of getting expected utility more than $2 - 1 = 1$. But (24) shows that s_B^* guarantees him 1. Therefore, B has no incentive to deviate from s_B^* . And the same argument holds for A , as well. Therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

Notice that this was already a new reasoning of being an equilibrium in a dynamic game. According to the knowledge of the author, \mathbf{s}^* does not satisfy the conditions of any other equilibrium concepts for dynamic games. For example, \mathbf{s}^* is not a Perfect Bayesian equilibrium, not even a subgame-perfect equilibrium for games with perfect information, because each player might miss the opportunity to completely utilize when the other player makes a bad move. As a simple counterexample, consider the following game. Player A has to choose his utility u_A from $[0, 1]$. Then player B , after observing u_A , has to choose his desired utility u'_B from $[0, 2]$. If $u_A + u'_B \leq 2$, then $u_B = u'_B$, otherwise $u_B = 0$. Then the strategy profile of choosing 1 by both players satisfies all (23), (24) and (25). But this is not subgame-perfect: if agent A chose e.g. $u_A = 0.7$, then the best choice for agent B would be $u'_B = 1.3$, but he chooses 1 instead. The only subgame-perfect equilibrium is that A chooses 1 and B chooses $2 - u_A$. But the utilities of the players are the same with both equilibria, and in this sense, we will show that quasi-dominant equilibrium is a stronger concept than Perfect Bayesian equilibrium.

Step 2. There is a set of players N playing a deterministic dynamic game with imperfect information. Suppose that there is a strategy profile $\mathbf{s}^* \in \mathcal{S}$ and constants C_i satisfying the

following.

$$\forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad u_i((s_i^*, \mathbf{s}_{-i})) \geq C_i \quad (26)$$

$$\forall \mathbf{s} \in \mathcal{S}: \quad \sum_{i \in N} u_i(\mathbf{s}) \leq \sum_{i \in N} C_i \quad (27)$$

Then \mathbf{s}^* is an equilibrium.

Justification. (26) implies that each player $i \in N$ can get utility at least C_i . (We say that C_i is *guaranteed* for i .) Each player $i \in N$ has no hope of getting more expected utility than the maximum possible total utility of all players minus the sum of the guaranteed utilities of the other players. Therefore, i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S}} \sum_{j \in N} u_j(\mathbf{s}) - \sum_{j \in N \setminus \{i\}} u_j(\mathbf{s}^*) \stackrel{(27)}{=} \sum_{j \in N} C_j - \sum_{j \in N \setminus \{i\}} C_j = C_i.$$

Consequently, each player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

Step 3. There is a set of players N playing a deterministic dynamic game with imperfect information. At the same initial time point, the players do simultaneous actions denoted by $a_i = \alpha(s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $a_i^* = \alpha(s_i^*)$ satisfying the following.

$$\forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad u_i((s_i^*, \mathbf{s}_{-i})) \geq f_i(\boldsymbol{\alpha}(\mathbf{s}_{-i})) \quad (28)$$

$$\forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall a_i \in \mathcal{A}_i: \quad \sum_{j \in N} u_j(\mathbf{s}) \leq f_i(\mathbf{a}_{-i}^*) + \sum_{j \in N \setminus \{i\}} f_j((a_i, \mathbf{a}_{-i-j}^*)) \quad (29)$$

Then \mathbf{s}^* is an equilibrium.

Justification. Each player $i \in N$ has no influence on \mathbf{a}_{-i} . Therefore, $f_i(\mathbf{a}_{-i})$ is independent of s_i , and (28) implies that i can get at least $f_i(\mathbf{a}_{-i})$ utility. Each player $i \in N$ has no hope of getting more expected utility than the maximum possible total utility of all players minus the sum of the guaranteed utilities of the other players. Therefore, if the players other than i take \mathbf{a}_{-i}^* , then i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S}} \sum_{j \in N} u_j(\mathbf{s}) - \inf_{a_i \in \mathcal{A}_i} \sum_{j \in N \setminus \{i\}} f_j((a_i, \mathbf{a}_{-i-j}^*)) \stackrel{(29)}{\leq} f_i(\mathbf{a}_{-i}^*).$$

Consequently, each player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

Step 4. There is a set of players N playing a dynamic *stochastic* game with imperfect information. At the same initial time point, the players do simultaneous actions denoted by $a_i = \alpha(s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $a_i^* = \alpha(s_i^*)$ satisfying the following.

$$\forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i((s_i^*, \mathbf{s}_{-i}))\right) \geq f_i(\boldsymbol{\alpha}(\mathbf{s}_{-i})) \quad (30)$$

$$\forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall a_i \in \mathcal{A}_i: \quad \sum_{j \in N} \mathbb{E}(u_j(\mathbf{s})) \leq f_i(\mathbf{a}_{-i}^*) + \sum_{j \in N \setminus \{i\}} f_j((a_i, \mathbf{a}_{-i-j}^*)) \quad (31)$$

Then \mathbf{s}^* is an equilibrium.

Justification. Each player $i \in N$ has no influence on \mathbf{a}_{-i} . Therefore, $f_i(\mathbf{a}_{-i})$ is independent of s_i , and (30) implies that i can get at least $f_i(\mathbf{a}_{-i})$ expected utility. Each player $i \in N$ has no hope of getting more expected utility than the maximum possible expected total utility of all players minus the sum of the guaranteed utilities of the other players. Therefore, if the players other than i take \mathbf{a}_{-i}^* , then i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S}} \sum_{j \in N} \mathbb{E}(u_j(\mathbf{s})) - \inf_{a_i \in \mathcal{A}_i} \sum_{j \in N \setminus \{i\}} f_j((a_i, \mathbf{a}_{-i-j}^*)) \stackrel{(31)}{\leq} f_i(\mathbf{a}_{-i}^*).$$

Consequently, each player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

Step 5. There is a set of players N playing a dynamic stochastic game with imperfect information. At the very beginning, nature assigns a *type* $\theta_i \in \Theta$ for each player $i \in N$, as a nondeterministic assignment with no assumed prior probabilities. Each player gets to know only his own type. Then at the same time point, the players do their first actions simultaneously, denoted by $a_i = \alpha(\theta_i, s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \Theta \times \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $\alpha^*(\theta_i) = \alpha(\theta_i, s_i^*)$, satisfying the following.

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i(\boldsymbol{\theta}, (s_i^*, \mathbf{s}_{-i}))\right) \geq f_i(\theta_i, \boldsymbol{\alpha}(\boldsymbol{\theta}_{-i}, \mathbf{s}_{-i})) \quad (32)$$

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall a_i \in \mathcal{A}_i:$$

$$\sum_{j \in N} \mathbb{E}(u_j(\boldsymbol{\theta}, \mathbf{s})) \leq f_i(\theta_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i})) + \sum_{j \in N \setminus \{i\}} f_j\left(\theta_j, (a_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i-j}))\right) \quad (33)$$

Then \mathbf{s}^* is an equilibrium.

Justification. In short, it is publicly known that whatever $\boldsymbol{\theta}$ is, if $\boldsymbol{\theta}$ was revealed, then \mathbf{s}^* would be an equilibrium as in Step 4. Therefore, we can rightfully say that \mathbf{s}^* is an equilibrium even without revealing $\boldsymbol{\theta}$.

A direct justification is the following. Each player $i \in N$ has no influence on $\boldsymbol{\theta}$ and \mathbf{a}_{-i} . Therefore, $f_i(\theta_i, \mathbf{a}_{-i})$ is independent of s_i , and (32) implies that i can get at least $f_i(\theta_i, \mathbf{a}_{-i})$ expected utility. Each player $i \in N$ has no hope of getting more expected utility than the maximum possible total expected utility of all players minus the sum of the guaranteed expected utilities of the other players. Therefore, if the players other than i take $\boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i})$, then i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S}} \sum_{j \in N} \mathbb{E}(u_j(\boldsymbol{\theta}, \mathbf{s})) - \inf_{a_i \in \mathcal{A}_i} \sum_{j \in N \setminus \{i\}} f_j(\theta_j, (a_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i-j})) \stackrel{(33)}{\leq} f_i(\theta_i, \boldsymbol{\alpha}^*(\boldsymbol{\theta}_{-i})).$$

Consequently, each risk-neutral player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

We call the equilibrium defined in Step 5 a **quasi-dominant equilibrium**.

8.1 The weak quasi-dominant equilibrium

We also introduce a slightly weaker version of the quasi-dominant equilibrium. We show the corresponding versions of Steps 4 and 5. The difference is that the function f may depend on a_i here. (Therefore, every quasi-dominant equilibrium is a weak quasi-dominant equilibrium.)

Step 4W. There is a set of players N playing a dynamic stochastic game with imperfect information. At the same initial time point, the players do simultaneous actions denoted by $a_i = \alpha(s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $a_i^* = \alpha(s_i^*)$ satisfying the following.

$$\forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i((s_i^*, \mathbf{s}_{-i}))\right) \geq f_i\left(\alpha(\mathbf{s}_{-i})\right) \quad (34)$$

$$\forall i \in N, \forall a_i \in \mathcal{A}_i, (\forall \mathbf{s} \in \mathcal{S} : \mathbf{a}(\mathbf{s}) = (a_i, \mathbf{a}_{-i}^*)):$$

$$\sum_{j \in N} \mathbb{E}(u_j(\mathbf{s})) \leq f_i(\mathbf{a}_{-i}^*) + \sum_{j \in N \setminus \{i\}} f_j((a_i, \mathbf{a}_{-i-j}^*)) \quad (35)$$

Then \mathbf{s}^* is an equilibrium.

Justification. Each player $i \in N$ has no influence on \mathbf{a}_{-i} . Therefore, if the others use \mathbf{a}_{-i}^* , then (34) implies that i can get at least $f_i(\mathbf{a}_{-i}^*)$ expected utility. Each player $i \in N$ has no hope of getting more expected utility than the maximum possible expected total utility of all players minus the sum of the guaranteed utilities of the other players. Therefore, if the players other than i take \mathbf{a}_{-i}^* , then i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S} : \alpha(\mathbf{s}_{-i}) = \mathbf{a}_{-i}^*} \left(\sum_{j \in N} u_j(\mathbf{s}) - \sum_{j \in N \setminus \{i\}} f_j((\alpha(s_i), \mathbf{a}_{-i-j}^*)) \right) \stackrel{(35)}{\leq} f_i(\mathbf{a}_{-i}^*).$$

Consequently, each player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

Step 5W. There is a set of players N playing a dynamic stochastic game with imperfect information. At the very beginning, nature assigns a *type* $\theta_i \in \Theta$ for each player $i \in N$, as a nondeterministic assignment with no assumed prior probabilities. Each player gets to know his own type. Then at the same time point, the players do their first actions simultaneously, denoted by $\alpha(\theta_i, s_i) \in \mathcal{A}_i$ for each player $i \in N$. Suppose that there are functions $f_i : \Theta \times \mathcal{A}_{-i} \rightarrow \mathbb{R}$ for all $i \in N$, and a strategy profile $\mathbf{s}^* \in \mathcal{S}$ with $\alpha^*(\theta_i) = \alpha(\theta_i, s_i^*)$, satisfying the following.

$$\forall \theta \in \Theta^N, \forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i(\theta, (s_i^*, \mathbf{s}_{-i}))\right) \geq f_i(\theta_i, \alpha(\theta_{-i}, \mathbf{s}_{-i})) \quad (36)$$

$$\forall \theta \in \Theta^N, \forall i \in N, (\forall \mathbf{s} \in \mathcal{S} \mid \alpha(\theta_{-i}, \mathbf{s}_{-i}) = \alpha^*(\theta_{-i})):$$

$$\sum_{j \in N} \mathbb{E}(u_j(\theta, \mathbf{s})) \leq f_i(\theta_i, \alpha^*(\theta_{-i})) + \sum_{j \in N \setminus \{i\}} f_j\left(\theta_j, (\alpha(\theta_i, s_i), \alpha^*(\theta_{-i-j}))\right) \quad (37)$$

Then \mathbf{s}^* is an equilibrium.

Justification. Each player $i \in N$ has no influence on θ and \mathbf{a}_{-i} . Therefore, if the other players use $\alpha^*(\theta_{-i})$, then (36) implies that i can get at least $f_i(\theta_i, \alpha^*(\theta_{-i}))$ expected utility. Each player $i \in N$ has no hope of getting more expected utility than the maximum possible total expected utility of all players minus the sum of the guaranteed expected utilities of the other players. Therefore, if the players other than i take $\alpha^*(\theta_{-i})$, then i has no hope of getting more expected utility than

$$\sup_{\mathbf{s} \in \mathcal{S}: \alpha(\theta_{-i}, \mathbf{s}_{-i}) = \alpha^*(\theta_{-i})} \left(\sum_{j \in N} \mathbb{E}(u_j(\theta, \mathbf{s})) - \sum_{j \in N \setminus \{i\}} f_j(\theta_j, (\alpha(\theta_i, s_i), \alpha^*(\theta_{-i-j}))) \right) \stackrel{(37)}{\leq} f_i(\theta_i, \alpha^*(\theta_{-i})).$$

Consequently, each risk-neutral player $i \in N$ has no incentive to deviate from s_i^* , and therefore, we can rightfully say that \mathbf{s}^* is an equilibrium. \square

We call the equilibrium defined in Step 5W a **weak quasi-dominant equilibrium**. Every quasi-dominant equilibrium is a weak quasi-dominant equilibrium because (33) implies (37).

8.2 The (weak) quasi-dominant equilibrium is a PBE-refinement

In this subsection, we add the necessary technical details to the justification of our equilibria in order to compare our equilibrium concepts with other solution concepts for dynamic games.

Bayesian belief system means the following. Nature has perfect information about the entire game and it sends special private messages to the players (beyond the other moves of nature specified in the rules of the game). The belief of each player is the history of all messages he received from all players including these special messages from nature. (Until this point, there was no loss of generality.) We assume that nature makes all of its moves according to a publicly known mixed strategy. This also means that if a game is defined with nondeterministic moves of nature (with no probabilities), then these moves become stochastic moves in the same game with Bayesian belief system.

The following proposition is the key to show that the weak quasi-dominant equilibrium is stronger than any reasonable equilibrium concept which (1) refines (Bayesian) Nash equilibrium and (2) which exists in all games (at least, in finite games). Of those, the most important concept is the Perfect Bayesian equilibrium.

Proposition 14. *Assume that a strategy profile \mathbf{s}^* is a weak quasi-dominant equilibrium in a game G . Denote by $G(i, \varepsilon)$ the game G with Bayesian belief system, and with the following modifications about the first parallel actions. Every player $j \in N \setminus \{i\}$ is forced to start with $\alpha^*(\theta_j)$. Player i is free to choose his first move except that deviating from $\alpha^*(\theta_i)$ has an extra cost of $\varepsilon > 0$ for him. Let $\tilde{\theta}_i$ denote the probabilistic variable of θ_i chosen by nature. Then all Nash equilibria \mathbf{s} of $G(i, \varepsilon)$ satisfy that $\alpha(\tilde{\theta}_i, s_i) = \alpha^*(\tilde{\theta}_i)$ with probability 1, and $\mathbb{E}(u_j(\tilde{\theta}, \mathbf{s}^*)) = \mathbb{E}(u_j(\tilde{\theta}, \mathbf{s}))$ for all players $j \in N$.*

Proof.

$$\begin{aligned} \mathbb{E}(u_i^G(\tilde{\theta}, \mathbf{s})) &\stackrel{(36)}{\leq} \mathbb{E}(u_i^G(\tilde{\theta}, \mathbf{s})) + \sum_{j \in N \setminus \{i\}} \mathbb{E}\left(u_j^G(\tilde{\theta}, (s_j^*, \mathbf{s}_{-j})) - f_j(\tilde{\theta}_j, \alpha(\tilde{\theta}_{-j}, \mathbf{s}_{-j}))\right) \\ &\leq \mathbb{E}(u_i^G(\tilde{\theta}, \mathbf{s})) + \sum_{j \in N \setminus \{i\}} \mathbb{E}\left(u_j^G(\tilde{\theta}, \mathbf{s}) - f_j(\tilde{\theta}_j, \alpha(\tilde{\theta}_{-j}, \mathbf{s}_{-j}))\right) \\ &= \sum_{j \in N} \mathbb{E}(u_j^G(\tilde{\theta}, \mathbf{s})) - \sum_{j \in N \setminus \{i\}} \mathbb{E}\left(f_j(\tilde{\theta}_j, \alpha(\tilde{\theta}_{-j}, \mathbf{s}_{-j}))\right) \end{aligned}$$

$$\begin{aligned}
&= \sum_{j \in N} \mathbb{E}(u_j^G(\tilde{\boldsymbol{\theta}}, \mathbf{s})) - \sum_{j \in N \setminus \{i\}} \mathbb{E}\left(f_j\left(\tilde{\theta}_j, (\alpha(\tilde{\theta}_i, s_i), \boldsymbol{\alpha}^*(\tilde{\boldsymbol{\theta}}_{-i-j}))\right)\right) \stackrel{(37)}{\leq} \mathbb{E}\left(f_i(\tilde{\theta}_i, \boldsymbol{\alpha}^*(\tilde{\boldsymbol{\theta}}_{-i}))\right) \\
&= \mathbb{E}\left(f_i(\tilde{\theta}_i, \boldsymbol{\alpha}(\tilde{\boldsymbol{\theta}}_{-i}, \mathbf{s}_{-i}))\right) \stackrel{(36)}{\leq} \mathbb{E}\left(u_i^G(\tilde{\boldsymbol{\theta}}, (s_i^*, \mathbf{s}_{-i}))\right) = \mathbb{E}\left(u_i^{G(i,\varepsilon)}(\tilde{\boldsymbol{\theta}}, (s_i^*, \mathbf{s}_{-i}))\right) \\
&\leq \mathbb{E}\left(u_i^{G(i,\varepsilon)}(\tilde{\boldsymbol{\theta}}, \mathbf{s})\right) = \mathbb{E}\left(u_i^G(\tilde{\boldsymbol{\theta}}, \mathbf{s})\right) - \varepsilon \cdot \mathbb{P}(\alpha(\tilde{\theta}_i, s_i) \neq \alpha^*(\tilde{\theta}_i)) \leq \mathbb{E}\left(u_i^G(\tilde{\boldsymbol{\theta}}, \mathbf{s})\right).
\end{aligned}$$

Therefore, all inequalities must hold with equality. For the last inequality, this means that

$$\mathbb{P}(\alpha(\tilde{\theta}_i, s_i) \neq \alpha^*(\tilde{\theta}_i)) = 0.$$

Notice that the entire calculation remains valid if we replace \mathbf{s} with \mathbf{s}^* . Consider now the first inequality of the calculation. Using that equality in the two cases,

$$\mathbb{E}\left(u_j^G(\tilde{\boldsymbol{\theta}}, \mathbf{s})\right) = \mathbb{E}\left(f_j(\tilde{\theta}_j, \boldsymbol{\alpha}(\tilde{\boldsymbol{\theta}}_{-j}, \mathbf{s}_{-j}))\right) = \mathbb{E}\left(f_j(\tilde{\theta}_j, \boldsymbol{\alpha}^*(\tilde{\boldsymbol{\theta}}_{-j}))\right) = \mathbb{E}\left(u_j^G(\tilde{\boldsymbol{\theta}}, \mathbf{s}^*)\right). \quad \square$$

Corollary 15. *If G is a finite game,⁹ then there exists a Perfect Bayesian equilibrium in each $G(i, 0)$ satisfying that $\alpha(\tilde{\theta}_i, s_i) = \alpha^*(\tilde{\theta}_i)$ with probability 1, and $\mathbb{E}(u_j(\tilde{\boldsymbol{\theta}}, \mathbf{s}^*)) = \mathbb{E}(u_j(\tilde{\boldsymbol{\theta}}, \mathbf{s}))$ for all players $j \in N$.*

Proof. G is finite, therefore, the set of mixed strategy profiles including beliefs is compact, and it contains at least one Perfect Bayesian equilibrium of each $G(i, \varepsilon)$ if $\varepsilon > 0$. Choose one for each $G(i, \frac{1}{k})$, where $k \in \mathbb{Z}$. These must have an accumulation point, and this point must be a Perfect Bayesian equilibrium of $G(i, 0)$. \square

We only need a very technical finishing step constructing a Perfect Bayesian equilibrium of G . We do not present it here, partially because it is very technical and not interesting and partially because there are multiple slightly different definitions of the Perfect Bayesian equilibrium in the literature which differ about the off-equilibrium paths, and these would require different constructions. But this is what we should conclude.

Corollary 16. *Assume that a strategy profile \mathbf{s}^* is a weak quasi-dominant equilibrium in a finite game G . Then in the game G with Bayesian belief system, there exists a Perfect Bayesian equilibrium \mathbf{s} satisfying that for all players $i \in N$, $\alpha(\tilde{\theta}_i, s_i) = \alpha^*(\tilde{\theta}_i)$ with probability 1, and $\mathbb{E}(u_i(\tilde{\boldsymbol{\theta}}, \mathbf{s}^*)) = \mathbb{E}(u_i(\tilde{\boldsymbol{\theta}}, \mathbf{s}))$, where $\tilde{\theta}_i$ denotes the probabilistic variable of θ_i chosen by nature.*

9 Proofs of quasi-dominant equilibria

9.1 The second-price mechanism

Theorem 1. *If the planner knows the capability tree of the principal, then the truthful strategy profile $\boldsymbol{\tau}$ is a quasi-dominant equilibrium under the second-price mechanism.*

Proof. We apply the equilibrium to our case. The first same-time actions of the players are identified with the proposals $\mathbf{a}_N = \hat{\boldsymbol{\theta}}_N$, where the principal has no other choice than $\hat{\theta}_0 = \theta_0$. Furthermore, $\mathbf{s}^* = \boldsymbol{\tau}$ with $\boldsymbol{\alpha}^*(\boldsymbol{\theta}) = \boldsymbol{\theta}$. What we need to prove is that there exist functions $f_i : \Theta^N \rightarrow \mathbb{R}$ for all player $i \in N$, so that the following inequalities hold.

⁹or a game with compact strategy spaces and with an upper semicontinuous utility function with respect to the topology on the executions induced by the product topology of the strategy spaces

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i(\boldsymbol{\theta}, (\boldsymbol{\tau}_i, \mathbf{s}_{-i}))\right) \geq f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i}(\boldsymbol{\theta}, \mathbf{s}_{-i})) \quad (38)$$

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall \hat{\theta}_i \in \Theta: \quad \mathbb{E}(u_N(\boldsymbol{\theta}, \mathbf{s})) \leq f_i(\boldsymbol{\theta}) + \sum_{j \in N \setminus \{i\}} f_j(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) \quad (39)$$

We show that the following functions f_i satisfy (38) and (39).

$$\forall i \in N^+: \quad f_i(\boldsymbol{\theta}) = V(\boldsymbol{\theta}) - V(\boldsymbol{\theta}_{-i}) \quad (40)$$

$$f_0(\boldsymbol{\theta}) = \sum_{i \in N^+} V(\boldsymbol{\theta}_{-i}) - (n-1)V(\boldsymbol{\theta}) \quad (41)$$

Proof of (38) for agents. $\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N^+, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}$:

$$\begin{aligned} \mathbb{E}\left(u_i^{2nd}(\boldsymbol{\theta}, (\boldsymbol{\tau}_i, \mathbf{s}_{-i}))\right) &\stackrel{(12)}{=} \mathbb{E}\left(u_i(\boldsymbol{\theta}, (\boldsymbol{\tau}_i, \mathbf{s}_{-i})) + v_i^+(\hat{\boldsymbol{\theta}})\right) = \mathbb{E}\left(u_i(\boldsymbol{\theta}, (\boldsymbol{\tau}_i, \mathbf{s}_{-i}))\right) + v_i^+(\theta_i, \hat{\theta}_i) \\ &\stackrel{(20)}{=} v_i^+(\theta_i, \hat{\theta}_i) \stackrel{(11)}{=} V((\theta_i, \hat{\boldsymbol{\theta}}_{-i})) - V(\hat{\boldsymbol{\theta}}_{-i}) \stackrel{(40)}{=} f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i}) \end{aligned}$$

Proof of (38) for the principal. $\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s}_{-0} \in \mathcal{S}_{-0}$:

$$\begin{aligned} \mathbb{E}(u_0^{2nd}(\boldsymbol{\theta}, \mathbf{s}_{-0})) &\stackrel{(13)}{=} \mathbb{E}(u_0(\boldsymbol{\theta}, \mathbf{s}_{-0})) - \sum_{i \in N^+} v_i^+(\hat{\boldsymbol{\theta}}) \stackrel{(18)(11)}{=} V(\hat{\boldsymbol{\theta}}) - \sum_{i \in N^+} (V(\hat{\boldsymbol{\theta}}) - V(\hat{\boldsymbol{\theta}}_{-i})) \\ &= \sum_{i \in N^+} V(\hat{\boldsymbol{\theta}}_{-i}) - (n-1)V(\hat{\boldsymbol{\theta}}) \stackrel{(41)}{=} f_0(\theta_0, \hat{\boldsymbol{\theta}}_{N^+}) \end{aligned}$$

Proof of (39). Notice that $\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N^+, \forall \hat{\theta}_i \in \Theta$:

$$\begin{aligned} \sum_{j \in N \setminus \{i\}} f_j(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) &= f_0(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) + \sum_{j \in N \setminus \{0, i\}} f_j(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) \stackrel{(41)(40)}{=} \sum_{j \in N \setminus \{0, i\}} V((\hat{\theta}_i, \boldsymbol{\theta}_{-i-j})) + \\ &+ V(\boldsymbol{\theta}_{-i}) - (n-1)V((\hat{\theta}_i, \boldsymbol{\theta}_{-i})) + \sum_{j \in N \setminus \{0, i\}} \left(V((\hat{\theta}_i, \boldsymbol{\theta}_{-i})) - V((\hat{\theta}_i, \boldsymbol{\theta}_{-i-j})) \right) = V(\boldsymbol{\theta}_{-i}). \end{aligned}$$

Therefore, $\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s} \in \mathcal{S}, \forall i \in N, \forall \hat{\theta}_i \in \Theta$:

$$\begin{aligned} f_i(\boldsymbol{\theta}) + \sum_{j \in N \setminus \{i\}} f_j(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) &\stackrel{(40)(41)}{=} (V(\boldsymbol{\theta}) - V(\boldsymbol{\theta}_{-i})) + \sum_{j \in N^+ \setminus \{i\}} (V(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) - V(\hat{\theta}_i, \boldsymbol{\theta}_{-i-j})) \\ &+ \sum_{j \in N^+ \setminus \{i\}} V(\hat{\theta}_i, \boldsymbol{\theta}_{-i-j}) + V(\boldsymbol{\theta}_{-i}) - (n-1)(V(\hat{\theta}_i, \boldsymbol{\theta}_{-i})) = V(\boldsymbol{\theta}) \stackrel{(5)}{\geq} \mathbb{E}(u_N(\boldsymbol{\theta}, \mathbf{s})). \quad \square \end{aligned}$$

9.2 The first-price mechanism under perfect competition

We say that there is **perfect competition** if it is common knowledge between the agents that the capability tree vector $\boldsymbol{\theta}$ is chosen from

$$\Theta_{\text{perf}}^N = \{\boldsymbol{\theta} \in \Theta^N \mid \forall i \in N^+: V(\boldsymbol{\theta}_{-i}) = V(\boldsymbol{\theta})\}. \quad (42)$$

Theorem 4. *The truthful strategy profile $\boldsymbol{\tau}$ under the first-price mechanism is a quasi-dominant equilibrium under perfect competition.*

Proof. We will use the following functions f_i :

$$\forall i \in \mathbf{N}^+ : \quad f_i(\boldsymbol{\theta}) = 0 \quad (43)$$

$$f_0(\boldsymbol{\theta}) = V(\boldsymbol{\theta}) \quad (44)$$

We will show that these functions satisfy (38) and (39) with “ $\forall \boldsymbol{\theta} \in \Theta_{\text{perf}}^N$ ” instead of “ $\forall \boldsymbol{\theta} \in \Theta^N$ ”. As the principal should also report her type, we also need to prove (39) for the principal, which is the following.

$$\forall \boldsymbol{\theta} \in \Theta_{\text{perf}}^N, \forall \hat{\boldsymbol{\theta}}_0 \in \Theta : \quad \sum_{i \in \mathbf{N}^+} f_i(\hat{\boldsymbol{\theta}}_0, \boldsymbol{\theta}_{\mathbf{N}^+}) \geq \sum_{i \in \mathbf{N}^+} f_i(\boldsymbol{\theta})$$

But in our case, both sides are 0.

Proof of (38) for agents. $\forall \boldsymbol{\theta} \in \Theta_{\text{perf}}^N, \forall i \in \mathbf{N}^+, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}$:

$$\mathbb{E}\left(u_i(\boldsymbol{\theta}, (\tau_i, \mathbf{s}_{-i}))\right) \stackrel{(20)}{=} 0 \stackrel{(43)}{=} f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i})$$

Proof of (38) for the principal. $\forall \boldsymbol{\theta} \in \Theta_{\text{perf}}^N, \forall \mathbf{s}_{-0} \in \mathcal{S}_{-0}$:

$$\mathbb{E}(u_0(\boldsymbol{\theta}, \mathbf{s}_{-0})) \stackrel{(18)}{=} V(\hat{\boldsymbol{\theta}}) \stackrel{(44)}{=} f_0(\theta_0, \hat{\boldsymbol{\theta}}_{\mathbf{N}^+})$$

Proof of (39). $\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s} \in \mathcal{S}, \forall i \in \mathbf{N}, \forall \hat{\boldsymbol{\theta}}_i \in \Theta$:

$$f_i(\boldsymbol{\theta}) + \sum_{j \in \mathbf{N} \setminus \{i\}} f_j(\hat{\boldsymbol{\theta}}_i, \boldsymbol{\theta}_{-i}) \stackrel{(43)(44)}{=} V(\hat{\boldsymbol{\theta}}_i, \boldsymbol{\theta}_{-i}) \geq V(\boldsymbol{\theta}_{-i}) \stackrel{(42)}{=} V(\boldsymbol{\theta}) \stackrel{(5)}{\geq} \mathbb{E}(u_{\mathbf{N}}(\boldsymbol{\theta}, \mathbf{s})). \quad \square$$

10 The first-price mechanism in general cases

We say that the planner **accepts** or **rejects** an agent depending on whether she asks i not to choose or to choose the do-nothing execution, respectively. The set of accepted agents is denoted by (winners) W .

For calculation purposes, we introduce the following notations. For any agent $i \in \mathbf{N}^+$, let Γ_{+i} denote the game we can obtain from Γ by excluding the option of the do-nothing execution from the capability tree θ_i of agent i . Similarly, Γ_{-i} is the game when $\theta_i = 0$, namely i is forced to do nothing. Finally, for a set of agents $S \subseteq \mathbf{N}^+$, Γ_S is the modified game when the agents in S do not have the option to do nothing, and all other agents are forced to do nothing.

We modify the first-price mechanism accordingly, namely, each agent with no do-nothing path has to submit a non-omissible capability tree (so these agent are automatically accepted). Let $M_i = \Gamma_{+i}/m_1$ and $M_{-i} = \Gamma_{-i}/m_1$ and $M_S = \Gamma_S/m_1$. We emphasize again that these are not interesting games but tools for calculation.

Notice that the proofs of Proposition 9 and Lemma 10 are valid in these games, as well. Namely, for any $M_* \in \{M_i, M_{-i}, M_S\}$,

$$\mathbb{E}(u_0^{M_*}) = V(M_*(\hat{\boldsymbol{\theta}})), \quad (45)$$

$$\mathbb{E}(u_i^{M_*}(\tau_i)) = 0.$$

Notice that

$$V(M) \stackrel{(5)}{=} \sup_{\mathbf{s} \in \mathcal{S}} \mathbb{E}(u_{\mathbf{N}}(\mathbf{s})) = \max \left(\sup_{\mathbf{s} \in \mathcal{S} \mid i \in W} \mathbb{E}(u_{\mathbf{N}}(\mathbf{s})), \sup_{\mathbf{s} \in \mathcal{S} \mid i \notin W} \mathbb{E}(u_{\mathbf{N}}(\mathbf{s})) \right) \stackrel{(5)}{=} \max(V(M_i), V(M_{-i})),$$

and the planner accepts i if $V(M_i(\hat{\theta})) > V(M_{-i}(\hat{\theta}))$, and rejects if $V(M_i(\hat{\theta})) < V(M_{-i}(\hat{\theta}))$. Notice that $V(M_{-i}(\hat{\theta})) = V(\hat{\theta}_{-i})$, because it does not matter whether i does not play at all or he takes part in the game but the planner rejects him for sure. Therefore,

$$V(\hat{\theta}) = \max \left(V(M_i(\hat{\theta})), V(\hat{\theta}_{-i}) \right). \quad (46)$$

Let the **signed value of a proposal** $\hat{\theta}_i$ be

$$v_i^\pm = v_i^\pm(\hat{\theta}) = v_i^\pm(\hat{\theta}_i) = V(M_i(\hat{\theta})) - V(M_{-i}(\hat{\theta})) = V(M_i(\hat{\theta})) - V(\hat{\theta}_{-i}). \quad (47)$$

Then

$$v_i^\pm \stackrel{(11)}{=} V(\hat{\theta}) - V(\hat{\theta}_{-i}) \stackrel{(46)}{=} \max \left(V(\hat{\theta}_{-i}), V(M_i(\hat{\theta})) \right) - V(\hat{\theta}_{-i}) \stackrel{(47)}{=} \max(0, v_i^\pm), \quad (48)$$

and $v_i^\pm > 0$ if and only if $i \in W$.

$V(M_i((\hat{\theta}_i - x, \hat{\theta}_{-i}))) = V(M_i(\hat{\theta})) - x$ because $M_i((\hat{\theta}_i - x, \hat{\theta}_{-i}))$ is the same game as $M_i(\hat{\theta})$ except that in the former game, agent i gets x less utility at the end. Thus,

$$v_i^\pm(\hat{\theta}_i - x) \stackrel{(47)}{=} V(M_i((\hat{\theta}_i - x, \hat{\theta}_{-i}))) - V(\hat{\theta}_{-i}) = V(M_i(\hat{\theta})) - x - V(\hat{\theta}_{-i}) \stackrel{(47)}{=} v_i^\pm(\hat{\theta}_i) - x.$$

Thus, for a particular agent, there is exactly one fair proposal with a given signed value.

Let the **strategy form** of a strategy s_i mean $\varphi(s_i) = \{s_i + x \mid x \in \mathbb{R}\}$. Let $\varphi(\tau_i) = \{\tau_i + x \mid x \in \mathbb{R}\}$ be called the **fair strategy form**.

Given an agent $i \in W$ and strategy form φ , i 's choice of strategy from φ has no direct effect on the other agents. Therefore, we assume that if an agent i uses another strategy from the same strategy form, and his proposal is accepted in both cases, then the execution of the game remains the same beyond the constant difference in the payment. This assumption automatically holds if the proposals are completely unobservable for the other agents.

Given θ and s_{-i} , we have that $E(u_i^{M_i}(s_i + x)) = E(u_i^{M_i}(s_i)) - x$. Furthermore, $v_i^\pm(\hat{\theta}_i + x) = v_i^\pm(\hat{\theta}_i) + x$. These imply that $E(u_i^{M_i}(s_i)) + v_i^\pm(\hat{\theta}_i)$ depends only on $\varphi(s_i)$. We call this sum the **value of the strategy form** $V_i(\varphi(s_i))$.

In each strategy form, there exists a strategy s_i , using a proposal $\hat{\theta}_i$, for which $\hat{\theta}_i + x$ would be accepted if $x > 0$ and rejected if $x < 0$. Then

$$E(u_i(s_i + x)) = \{0 \text{ if } x < 0; \text{ and } V_i(\varphi(s_i)) - x \text{ if } x > 0\}.$$

Let a **fair-looking strategy** of an agent $i \in N^+$ mean a strategy by which i sends a proposal $\hat{\theta}_i$, and then he behaves as follows. The reported outcome of the chance event has the same probability distribution as in $\hat{\theta}_i$ which is independent of the preceding actions of all other players, and in the end, he provides the consequence corresponding to the communication. In other words, he behaves as an agent with capability tree $\hat{\theta}_i$ who uses truthful strategy.

Clearly, the truthful strategy is the fair strategy with profit 0, and all fair strategies are fair-looking strategies. The primary examples of fair-looking strategies is when the agent sends the proposal of his real capability tree but with valuation decreased by different amounts at different executions, and then he plays truthfully and obediently.

Similarly to Lemma 10, the following lemma shows that an agent i with fair strategy $\tau_i - x$ has a guaranteed expected utility x in the case of acceptance.

Lemma 17. *If $i \in W$, then*

$$E(u_i(s_i = \tau_i - x)) = x. \quad (49)$$

Proof. Notice that if $s_i = \tau_i - x$, then $\hat{v}_i = v_i - x$. Therefore,

$$\mathbb{E}(u_i) \stackrel{(3)}{=} \mathbb{E}(v_i + p_i) \stackrel{(9)}{=} \mathbb{E}\left(v_i - \hat{v}_i + \sum_{\chi \in X_i} \delta(\chi)\right) \stackrel{(17)}{=} \mathbb{E}(v_i - \hat{v}_i) = x. \quad \square$$

The following theorem shows that if all agents use fair strategy, then after the choice of W , the mechanism works efficiently. This implies that the planner chooses the set W for which the expected total utility *minus the sum of the demanded profits of the agents in W* is the largest possible; and the choice of W is the only inefficient decision throughout the game. Let $W_0 = W \cup 0$. We use the notation M_S for $S \subseteq N^+$ expressing that S must be accepted, defined in the beginning of Section 10.

Theorem 18. *If $\forall i \in W : s_i = \tau_i - x_i$, then*

$$\mathbb{E}(u_N) = V(M_W) \quad (50)$$

Proof. Notice that the only difference between $M_W(\boldsymbol{\theta})$ and $M_W(\hat{\boldsymbol{\theta}})$ is that each agent $i \in W$ gets x_i lower utility in the latter game. Also, $M_W(\boldsymbol{\theta}'_N)$ and $M_W(\boldsymbol{\theta}'_{W_0})$ are always the same except that $M_W(\boldsymbol{\theta}'_N)$ includes a subset of agents $N^+ \setminus W$ who must be rejected.

Hence,

$$u_0^{M_W} \stackrel{(45)}{=} V(M_W(\hat{\boldsymbol{\theta}}_{W_0})) = V(M_W) - \sum_{i \in W} x_i, \quad (51)$$

therefore, in M_W ,

$$\mathbb{E}(u_N) = \mathbb{E}(u_0) + \sum_{i \in W} \mathbb{E}(u_i) \stackrel{(51)(49)}{=} \left(V(M_W) - \sum_{i \in W} x_i \right) + \sum_{i \in W} x_i = V(M_W). \quad \square$$

Theorem 19. *Consider an arbitrary agent $i \in N^+$. Suppose that every other agent $j \in W \setminus \{i\}$ uses a fair-looking strategy s_j . Then the fair strategies of i maximize the value of the strategy form*

$$v_i^\pm \left(\hat{\theta}((s_i, \mathbf{s}_{-i})) \right) + \mathbb{E} \left(u_i^{M_i}((s_i, \mathbf{s}_{-i})) \right).$$

Proof. Notice that whether an agent j has a capability tree θ_j and he uses a fair-looking strategy s_j , or his capability tree is $\hat{\theta}_j(\theta_j, s_j)$ and he uses truthful strategy, these are the same from all the other players' point of view. Therefore, from the point of view of i , the only difference between $M(\boldsymbol{\theta})$ with fair-looking strategies \mathbf{s}_{-i} and $M' = M((\theta_i + x, \hat{\boldsymbol{\theta}}_{-i}(\mathbf{s}_{-i})))$ with truthful strategies $\mathbf{s}_{-i}^{M'} = \boldsymbol{\tau}_{-i}^{M'}$ is that the utility of i is lower by x and the signed value of his proposal is higher by x in the latter case. Thus, after this transformation of the game M to M' with an x large enough, we only need to prove the following. If $\mathbf{s}_{-i} = \boldsymbol{\tau}_{-i}$ and any s_i with proposal $\hat{\theta}_i$ satisfies $v_i^\pm(\boldsymbol{\theta}) = v_i^\pm((\hat{\theta}_i, \boldsymbol{\theta}_{-i})) > 0$, then $\mathbb{E}(u_i^{M_i}(s_i, \boldsymbol{\tau}_{-i})) \leq \mathbb{E}(u_i^{M_i}(\boldsymbol{\tau}))$.

For a large x , $i \in W$, therefore,

$$\begin{aligned} & \mathbb{E} \left(u_i^{M_i}((s_i, \boldsymbol{\tau}_{-i})) \right) = \mathbb{E} \left(u_i((s_i, \boldsymbol{\tau}_{-i})) \right) \\ & = \mathbb{E} \left(u_N((s_i, \boldsymbol{\tau}_{-i})) \right) - \mathbb{E} \left(u_0((s_i, \boldsymbol{\tau}_{-i})) \right) - \sum_{j \in N^+ \setminus \{i\}} \mathbb{E} \left(u_j((s_i, \boldsymbol{\tau}_{-i})) \right) \end{aligned}$$

$$\stackrel{(6)(18)(20)}{\leq} V(M) - V \left(\hat{\theta}((s_i, \boldsymbol{\tau}_{-i})) \right) = V(M) - V \left(M_i(\hat{\boldsymbol{\theta}}((s_i, \boldsymbol{\tau}_{-i}))) \right) \stackrel{(47)}{=} V(M) - V(\hat{\boldsymbol{\theta}}_{-i}) - v_i^\pm$$

and Theorem 11 shows that equality holds if $s_i = \tau_i$. \square

$\mathbf{E}(u_i(s_i)) \leq V_i(\varphi(s_i))$, therefore, the strategy form of $i \in \mathbf{N}^+$ with the highest value gains him the highest potential for his expected utility. Theorem 19 implies that the fair strategy form has the highest value. Moreover,

$$\mathbf{E}(u_i(\tau_i - x)) = \{x \text{ if } V_i(\tau_i) > x; \text{ and } 0 \text{ otherwise}\},$$

which is quite an efficient way of the exploitation of this potential. Of course, this only shows that the fair strategy with appropriate profit is approximately the best for an agent, under competition.

The following arguments show that under practically reasonable approximations or assumptions, the agents have very weak or no incentives in using not fair strategies.

We consider the Project Management Model Γ with Bayesian belief system (see Section 8.2). Now θ is a stochastic variable. Let us fix an agent $i \in \mathbf{N}^+$ and the mixed strategies of others \mathbf{s}_{-i} . This implies that the probability distribution of $\hat{\theta}_{-i}$ is also fixed. Let \mathbf{P}_i and \mathbf{E}_i denote the probability and expectation given \mathbf{s}_{-i} and the belief of $i \in \mathbf{N}^+$ at t_0 . Consider a strategy s_i and denote by $\hat{\theta}_i$ the proposal he would send by s_i . Let (value) $V^\pm = v_i^\pm(\theta_i, \hat{\theta}_{-i})$, (difference) $D = V^\pm - v_i^\pm(\hat{\theta}) = V(\theta_i, \hat{\theta}_{-i}) - V(\hat{\theta})$ and $e = e(\hat{\theta}_i) = \mathbf{E}_i(D \mid D < V^\pm)$.

In practice, V^\pm and D are almost independent and both have ‘‘natural’’ distributions, therefore, $\mathbf{P}_i(e < V^\pm) = \mathbf{P}(\mathbf{E}_i(D \mid D < V^\pm) < V^\pm)$ is usually not smaller than $\mathbf{P}_i(D < V^\pm)$. This observation shows the importance of the following theorem.

Theorem 20. *If \mathbf{s}_{-i} consists only of fair-looking strategies and $\mathbf{P}_i(\mathbf{E}_i(D \mid D < V^\pm) < V^\pm) \geq \mathbf{P}_i(D < V^\pm)$ holds for all s_i , then a fair strategy of i provides him the highest expected utility.*

Proof. Let $\bar{u}(s_i) = \mathbf{E}_i(u_i(s_i))$ and let x be the number by which $\bar{u}(\tau_i + x)$ is the largest possible. What we need to prove is that $\forall s_i: \bar{u}(\tau_i + x) \geq \bar{u}(s_i)$.

Theoretically, let us allow for $i \in \mathbf{N}^+$ to submit his fair proposal with the signed value $v_i^\pm(\hat{\theta}'_i, \hat{\theta}_{-i})$, namely, submitting $\theta_i + v_i^\pm(\theta_i, \hat{\theta}_{-i}) - v_i^\pm(\hat{\theta}'_i, \hat{\theta}_{-i})$, for an arbitrary θ'_i chosen by i . Normally, it is an invalid action, because i does not know the values of his possible proposals at this time, but for calculation, we allow him this, and we denote the fair strategy but with this proposal by $\text{fair}_i(\theta'_i)$.

$\hat{\theta}_i$ is accepted if and only if $v_i^\pm(\hat{\theta}) > 0$, or equivalently, $D < V^\pm$. By the equation

$$\bar{u}(s_i) = \mathbf{P}_i(i \in W) \cdot \mathbf{E}_i(u_i(s_i) \mid i \in W),$$

we get $\bar{u}(\tau_i - e) \stackrel{(49)}{=} \mathbf{P}_i(e < V^\pm)e$, and $\bar{u}(\text{fair}_i(\hat{\theta}_i)) = \mathbf{P}_i(D < V^\pm)e$, whence we can simply get that

$$\begin{aligned} & \bar{u}(\tau_i - x) - \bar{u}(s_i) = \\ & = \left(\bar{u}(\tau_i - x) - \bar{u}(\tau_i - e) \right) + \left(\mathbf{P}_i(e < V^\pm) - \mathbf{P}_i(D < V^\pm) \right) e + \left(\bar{u}(\text{fair}_i(\hat{\theta}_i)) - \bar{u}(s_i) \right). \end{aligned}$$

- $\bar{u}(\tau_i - x) - \bar{u}(\tau_i - e) \geq 0$ by the definition of x .
- If $e \leq 0$, then $\bar{u}(s_i) \leq 0 = \bar{u}(\tau_i) \leq \bar{u}(\tau_i - x)$, therefore, s_i cannot be a better strategy. Assume that $e > 0$. In this case, because of the assumption in the theorem,

$$\left(\mathbf{P}_i(e < V^\pm) - \mathbf{P}_i(D < V^\pm) \right) e \geq 0.$$

- Theorem 18 implies that $\bar{u}(\text{fair}_i(\hat{\theta}_i)) - \bar{u}(s_i) \geq 0$.

To sum up, $\bar{u}(\tau_i - x) - \bar{u}(s_i) \geq 0$, which proves the theorem. \square

10.1 Negative results

First, we show what kind of deviations from fair strategies we should expect. Consider a case when an agent $i \in N^+$ will be asked to work on either a larger or a smaller task, and this choice depends only on the proposals of others. Then i should demand more profit (payment beyond his costs) in the former case and less profit in the latter case, due to the different competition settings in the two cases. Say, if the larger task has the very same structure, but it has double valuations for everyone, then it is easy to check that the agents should demand exactly double profit in this case. And this is not a fair proposal.

The opinion of the author is that a rational agent would typically use a strategy very similar to a fair strategy, mainly with deviations like in the previous case. This strategy would still provide approximate efficiency. However, we describe a situation when an agent would use a highly non-fair-looking strategy, under a very extreme belief system.

Consider an agent i working a subtask. Suppose that agent i knows that his capability tree is clearly the best one for the task, and agent j is clearly the second best. Suppose that i knows very well the capability tree and belief of j , but j believes that with probability close to 1, i knows almost nothing about the capability tree of j . In this case, it can easily happen that i can predict well the proposal of j . If i does not know the capability trees of agents working on other tasks, then i should submit the proposal of agent j but claiming slightly smaller payment, in order to very slightly underbid j . This can provide i with almost the maximum possible expected utility $V(\theta_i, \hat{\theta}_{-i}) - V(\hat{\theta}_{-i})$. Further techniques related to this question are shown in Appendices E.1 and E.2.

We note that the belief structure of the example above is exceptional in the sense that under common prior assumption, this must have very low probability, and even this unlikely event causes moderate inefficiency.

11 First-price mechanism with commonly known capability trees

In this section, we define the game so that the principal reports her capability tree, the planner determines the best set of agents, and the rejected agents must do nothing (rather than just incentivized to do so).

Assume that the capability trees of all players are common knowledge between them. We emphasize that the planner does not know the capability trees and the entire executions might be completely hidden.

In order to avoid some unimportant technical difficulties, somewhat imprecisely, we introduce a positive infinitesimal amount of money ε , and we will consider it at the decision about which agents to accept, but we will neglect it in all other aspects. (In fact, we are considering limits of ε -equilibria with $\varepsilon \rightarrow 0$.) Furthermore, for the sake of simplicity, we assume that there is no tie, namely, $W = \arg \max_{S \subseteq N^+} V(M_W)$, or equivalently, if $\hat{\theta} = \theta$, then $W = \{i \in N^+ : v_i^+(\theta) > 0\}$.

First, we show an example why we should expect multiple equilibria when the capability trees are common knowledge. Assume that there are 3 agents. Agent A could complete one half of the project with a cost of 4, agent B could complete the other half also with a cost of 4, and agent C could complete the entire project with a cost of 10. Now, for any $x \in (4, 6)$, if agents A , B and C report a cost of x (or $x - \varepsilon$), and $10 - x$ (or $10 - x - \varepsilon$) and 10, respectively then this is a (2ε) -equilibrium.

Consider the strategy profile $\boldsymbol{\tau} - \mathbf{c} = (\tau_i - c_i)_{i \in N}$, where the vector c_i satisfies

$$c_i = \begin{cases} 0 & \text{if } i = 0 \text{ or } v_i^+(\boldsymbol{\theta}) = 0, \\ v_i^+(\theta_i, \hat{\boldsymbol{\theta}}_{-i}) - \varepsilon > 0 & \text{if } i \neq 0 \text{ and } v_i^+(\boldsymbol{\theta}) > 0. \end{cases} \quad (52)$$

Or equivalently, the choice of \mathbf{c} satisfies that if $v_i^+(\boldsymbol{\theta}) > 0$, then $v_i^+(\hat{\boldsymbol{\theta}}) = \varepsilon$ and $c_i > 0$, otherwise $c_i = 0$. This expresses that all players use a fair proposal, the best set of agents overbid their real costs by ε less than what would cause a tie, while others report their true capability trees.

This way, $u_0 = V(\boldsymbol{\theta} - \mathbf{c})$, and for $i \in N^+$, $u_i = c_i$, and Theorem 18 implies that they maximize the expected total utility, therefore,

$$V(\boldsymbol{\theta}) = V(\boldsymbol{\theta} - \mathbf{c}) + \sum_{i \in N} c_i. \quad (53)$$

The vector \mathbf{c} is an ε -Pareto-efficient core solution of the superadditive coalitional game for the agents defined by V . The existence of such a vector is a very simple and well-known fact in cooperative game theory. We can find one by the following procedure. We start from $\mathbf{c} \equiv 0$, and in an arbitrary ordering, we keep increasing each value c_i , $i \in W$ as long as $v_i^+(\boldsymbol{\theta} - \mathbf{c}) > \varepsilon$. The resulting vector will satisfy (52). It is also easy to see that each $c_i \leq v_i^+(\boldsymbol{\theta})$.¹⁰

$\boldsymbol{\tau} - \mathbf{c}$ maximizes the expected total utility and provides at most as much expected utility for the agents as under the second-price mechanism. Therefore, by exclusion, *the principal gets at least as much utility here as under the second-price mechanism.*

Now we show that $\boldsymbol{\tau} - \mathbf{c}$ is also a quasi-dominant equilibrium, which completes the proof of Theorem 5.

Theorem 21 (Specified version of Theorem 5). *If the capability trees $\boldsymbol{\theta}$ are commonly known, then $\boldsymbol{\tau} - \mathbf{c}$ is a weak quasi-dominant equilibrium under the first-price mechanism (in the limit $\varepsilon \rightarrow 0$)¹¹.*

Proof. We need to find functions f_i which satisfy (36) and (37), which equations are the following in our case.

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}: \quad \mathbb{E}\left(u_i(\boldsymbol{\theta}, (\tau_i - c_i, \mathbf{s}_{-i}))\right) \geq f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i}(\boldsymbol{\theta}, \mathbf{s}_{-i})) \quad (54)$$

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N, (\forall \mathbf{s} \in \mathcal{S} \mid \hat{\boldsymbol{\theta}}_{-i}(\boldsymbol{\theta}, \mathbf{s}_{-i}) = (\boldsymbol{\theta} - \mathbf{c})_{-i}):$$

$$\mathbb{E}(u_N(\boldsymbol{\theta}, \mathbf{s})) \leq f_i(\theta_i, (\boldsymbol{\theta} - \mathbf{c})_{-i}) + \sum_{j \in N \setminus \{i\}} f_j\left(\theta_j, (\hat{\boldsymbol{\theta}}_i, (\boldsymbol{\theta} - \mathbf{c})_{-i-j})\right) \quad (55)$$

We prove that the following functions f_i satisfy (54) and (55). $I(\text{true}) = 1$ and $I(\text{false}) = 0$.

$$\forall i \in N^+: \quad f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i}) = I(i \in W \mid \hat{\boldsymbol{\theta}}_{-i}; \hat{\theta}_i = \theta_i - c_i) \cdot c_i \quad (56)$$

$$f_0(\theta_0, \hat{\boldsymbol{\theta}}_{N^+}) = V(\theta_0, \hat{\boldsymbol{\theta}}_{N^+}) \quad (57)$$

Proof of (54) for the agents. $\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{s}_{-i} \in \mathcal{S}_{-i}$:

$$\mathbb{E}\left(u_i(\boldsymbol{\theta}, (\tau_i - c_i, \mathbf{s}_{-i}))\right) \stackrel{(49)}{=} I(i \in W \mid \boldsymbol{\theta}; (\tau_i - c_i, \mathbf{s}_{-i})) \cdot c_i \stackrel{(56)}{=} f_i(\theta_i, \hat{\boldsymbol{\theta}}_{-i}(\boldsymbol{\theta}, \mathbf{s}_{-i}))$$

¹⁰The payoff of each player i by any core solution of a superadditive coalitional game is at most the marginal contribution of i . Otherwise the other players would get more payoff in total by excluding i from the coalition.

¹¹This is just the same technical issue as what we have even in single-item auctions: the bidder with the highest valuation should overbid the second highest valuation by some $\varepsilon > 0$.

Proof of (54) for the principal. $\forall \boldsymbol{\theta} \in \Theta^N$, $\forall \mathbf{s}_{N^+} \in \mathcal{S}_{N^+}$:

$$\mathbb{E}\left(u_0(\boldsymbol{\theta}, (\tau_0, \mathbf{s}_{N^+}))\right) \stackrel{(18)}{=} V(\theta_0, \hat{\boldsymbol{\theta}}_{N^+}(\boldsymbol{\theta}, \mathbf{s}_{N^+})) \stackrel{(57)}{=} f_0(\theta_0, \hat{\boldsymbol{\theta}}_{N^+}(\boldsymbol{\theta}, \mathbf{s}_{N^+}))$$

Proof of (55) for the agents. Let $W^* = W(\hat{\theta}_i, (\boldsymbol{\theta} - \mathbf{c})_{-i})$.

$\forall \boldsymbol{\theta} \in \Theta^N$, $\forall i \in N^+$, $(\forall \mathbf{s} \in \mathcal{S} \mid \hat{\boldsymbol{\theta}}_{-i}(\boldsymbol{\theta}, \mathbf{s}_{-i}) = (\boldsymbol{\theta} - \mathbf{c})_{-i})$:

$$\begin{aligned} & f_i(\theta_i, (\boldsymbol{\theta} - \mathbf{c})_{-i}) + \sum_{j \in N \setminus \{i\}} f_j(\theta_j, (\hat{\theta}_i, (\boldsymbol{\theta} - \mathbf{c})_{-i-j})) \\ & \stackrel{(56)(57)(52)}{=} V\left((\hat{\theta}_i, (\boldsymbol{\theta} - \mathbf{c})_{-i})\right) + \sum_{j \in W^* \cup \{i\}} c_j \geq V((\boldsymbol{\theta} - \mathbf{c})_{-i}) + \sum_{j \in W^*} c_j \\ & \stackrel{(52)}{=} V(\boldsymbol{\theta} - \mathbf{c}) + \sum_{i \in W^*} c_i \stackrel{(53)}{\geq} V\left((\boldsymbol{\theta}_{W^*}, (\boldsymbol{\theta} - \mathbf{c})_{N \setminus W^*})\right) \geq V(\boldsymbol{\theta}_{W^* \cup \{0\}}) \geq \mathbb{E}(u_N(\boldsymbol{\theta}, \mathbf{s})) \quad (58) \end{aligned}$$

Proof of (55) for the principal. Let $W^* = W(\hat{\theta}_0, (\boldsymbol{\theta} - \mathbf{c})_{N^+})$.

$\forall \boldsymbol{\theta} \in \Theta^N$, $(\forall \mathbf{s} \in \mathcal{S} \mid \hat{\boldsymbol{\theta}}_{N^+}(\boldsymbol{\theta}, \mathbf{s}_{N^+}) = (\boldsymbol{\theta} - \mathbf{c})_{N^+})$:

$$f_0(\theta_0, (\boldsymbol{\theta} - \mathbf{c})_{N^+}) + \sum_{i \in N^+} f_i(\theta_i, (\hat{\theta}_0, (\boldsymbol{\theta} - \mathbf{c})_{N^+})) \stackrel{(57)(56)(52)}{=} V(\boldsymbol{\theta} - \mathbf{c}) + \sum_{i \in W^*} c_i \stackrel{(58)}{=} \mathbb{E}(u_N(\boldsymbol{\theta}, \mathbf{s})) \quad \square$$

Acknowledgements

I am thankful to **Péter Eső** for his suggestions for revision of the paper, and to **Miklós Péter Pintér** for the suggestions about the earlier version (MSc Thesis) of the paper. I am also very grateful for the help of Mihály Bárász, Gergely Csapó, Katalin Friedl, László Lovász, Éva Tardos, Alexander Teytelboym, László Végh and many other people.

References

- [1] S. Athey and I. Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.
- [2] D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- [3] R. Cavallo, D.C. Parkes, and S. Singh. Optimal coordinated planning amongst self-interested agents with private state. *arXiv preprint arXiv:1206.6820*, 2012.
- [4] E.J. Friedman and D.C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 240–241. ACM, 2003.
- [5] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [6] D.C. Parkes and S.P. Singh. An mdp-based approach to online mechanism design. In *Advances in neural information processing systems*, 2003.
- [7] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Appendix

A Negative results and counterexamples

A.1 Dominant strategy equilibrium is not possible

Consider a project consisting of two tasks, carried out by two players: A and B . First, A completes the first task with two possible qualities: high or low. Then B completes the second task, using two different technologies in the two cases.

Assume that B can make some preparation for using either or both technologies before A starts. Preparation to each case costs 1, but it saves him 3 cost in the corresponding case. Which preparation(s) B makes remains hidden throughout the game, therefore, this decision has no effect on the transfer to him. The probability of the high quality depends on the non-observable effort level of A . Therefore, the optimal decision of B depends on the hidden strategy of A . This proves the impossibility of a dominant strategy equilibrium.

A.2 A counterexample about belief-independence

Assume that there is a game with perfect information, and there is an equilibrium \mathbf{s} with which each player uses only a part of his information. Consider now the same game with an imperfect information setup satisfying that all players still have the information which is required to follow the same strategy in \mathbf{s} . There is a well-known principle saying that under some assumptions, \mathbf{s} must be an equilibrium in this new game. The following example shows that this principle is not valid in dynamic stochastic environments including our Project Management Model. This is the reason why we use different techniques.

There are two players: A and B . At the beginning A decides whether to participate in the game for a cost of 1. If not, then the game ends with utilities $(0, 0)$. If yes, then B can choose not playing, resulting in utilities $(-1, 0)$, or continuing. Then B should make a decision between two options called red and blue, and then A makes a guess for the decision. If the guess is right, then A gets 5, and B gets -1 , which are utilities $(4, -1)$ in total, but if wrong, then B gets 5 and A gets -1 , which are utilities $(-2, 5)$ in total.

In the game with perfect information, B would not play, because A would see his choice, and would make the right guess. Therefore A is better off not participating in the game. To sum up, non-participation strategy for both players (with arbitrary decision and guess) is the (essentially) unique equilibrium. With this equilibrium, A does not use the information on B 's choice of color. Therefore, the principle would say that this must be an equilibrium under imperfect information even if A does not observe the decision of B .

Consider now the same game when A cannot observe the decision of B . In this case, B would be happy to participate, because this would provide him with utility $\frac{5-1}{2} = 2$ by making a uniform random choice. Therefore, A should participate as well, and make a random guess, which provides her with utility $\frac{5-1}{2} - 1 = 1$. This contradicts with the principle.

A.3 Collusion under the second-price mechanism

Consider a case when two agents can send proposals defined in such a way that their individual work is useless without each other. For example, agent 1 has a pink dummy spaceship in neo-Hawaiian style, and he offers to bring it; and agent 2 offers to complete the task provided that he gets a pink dummy spaceship in neo-Hawaiian style, as a tool for his work. And, of course,

nobody else has or requires such a toy. Assume that these proposals have positive values. If agent 2 reports a valuation increased by a constant x (reports x lower cost for all executions) except for the do-nothing execution, then the value of both proposals would increase by x . Consequently, they would get $2x$ more second-price compensation in total, which means x more total utility for them. Using this trick, these players can get as much payment as they want.

B Comparison with existing models

First, we show a somewhat lossy translation of our Project Management Model to the language used by Athey and Segal (2013) [1]. This discretized version is difficult to interpret correctly and we omit some minor details here, but still this version can be useful for a comparison of our second-price mechanism with previous results. We strongly suggest reading this section only after understanding the model (Sections 2, 4).

We have a set $N = \{1, 2, \dots, n\} = N^+ \cup N^-$ of agents, $N^+ = \{1, 2, \dots, k\}$, $N^- = \{k + 1, k + 2, \dots, n\}$, $k \geq 1$, and there is a countable number of periods $t \in \mathbb{N} = \{0, 1, \dots\}$.

Period 0 consists of the following steps.

- Nature chooses the type vector $\theta_0 \in \Theta^N$ arbitrarily, with no prior distribution.
- Each agent $i \in N$ privately observes his type $\theta_0^i \in \Theta$.
- The planner observes $\theta_0^{N^+}$.
- Each agent $i \in N^-$ has an option not to join the game.

Each period $t \in \{1, 2, \dots\}$ consists of the following steps.

- Each agent $i \in N$ makes a private decision $x_t^i \in X$.
- The planner makes a public decision $x_t^0 \in X$.
- Each agent $i \in N$ privately observes his realized private state $\theta_t^i \in \Theta$.
- All agents observe a realized public state $\theta_t^0 \in \Theta$.
- A public consequence vector $\mathbf{c}_t \in C^N$ is announced. Each c_t^i is a function of x_t^i, x_t^0, θ_t^i and θ_t^0 . If i did not join to the game, then $c_t^i = \emptyset$.
- Each agent $i \in N$ is given a transfer $y_t^i \in \mathbb{R}$ with $\sum_{i \in N} y_t^i = 0$. If agent $i \in N^-$ did not join, then $y_t^i = 0$.

The utility of each agent $i \in N$ is given as a function of the sequences of his private states $(\theta_t^i)_{t=0}^\infty$, the public states $(\theta_t^0)_{t=0}^\infty$, his private decisions $(x_t^i)_{t=1}^\infty$, the public decisions $(x_t^0)_{t=1}^\infty$, the consequences $(\mathbf{c}_t)_{t=1}^\infty$ and monetary transfers $(y_t^i)_{t=1}^\infty$, as follows:

$$\text{utility}(i) = \sum_{t=1}^{\infty} \delta^t \left(v^i(x_t^i, x_t^0, \theta_t^i, \theta_t^0, \mathbf{c}_t) + y_t^i \right),$$

where $\delta \in (0, 1)$ is a discount factor.

The distribution of subsequent states of each agent $i \in N$ is governed by transition probability measure $\mu^i: X^2 \times \Theta^2 \times C^N \rightarrow \Delta(\Theta)$. Specifically, for any period $t \geq 0$, the next period

state of $i \in N$ is a random variable $\tilde{\theta}_{t+1}^i$ distributed according to the probability measure $\mu^i(x_t^i, x_t^0, \theta_t^i, \theta_t^0, \mathbf{c}_t) \in \Delta(\Theta)$. Similarly, $\tilde{\theta}_{t+1}^0$ is distributed according to $\mu^0(x_t^0, \theta_t^0, \mathbf{c}_t) \in \Delta(\Theta)$.

We assume that for all $t \in \{1, 2, \dots\}$, the beliefs of the players in $N \setminus \{i\}$ just after period t (when i can make a report about θ_t^i) are independent from θ_t^i conditional on all beliefs of all agents before the state was chosen from the distribution.

We design a mechanism that makes the public decisions and determines monetary transfers depending on the information of the planner including the history of public reports of the agents. We show that this implements the efficient strategy profile in a quasi-dominant equilibrium.

We note that using consequences or public decisions are equivalent. On one hand, a consequence is equivalent to a public decision where deviation is punished by the monetary transfer rule. On the other hand, having a public decision is equivalent to having a private decision with a consequence rule which reveals the decision.

B.1 Comparison to Athey–Segal

Similar questions in dynamic mechanism design were analyzed by Athey and Segal (2013) [1]. Here we consider a comparison between their model with the “balanced team mechanism” and our Project Management Model with our “second-price mechanism”.

In their model, the initial state $\theta_0 \in \Theta$ is assumed to be publicly known. In contrast, in our Project Management Model all but at least one agent can have a private initial state chosen by nature, with no prior, meaning that our solution concept will be ex post with respect to the initial types. But we assume that those agents have the option of doing nothing and if they do so or do anything providing no consequence, then they are indifferent about the execution of the game.

Their model uses a countable number of time periods. During each period, the players may receive private signals which they should report at the end of the period. Therefore, in that paper they had to use ex post incentive compatibility with respect to the same-period signals of the other players. E.g. they had to consider both of the possibilities that an agent i misreports depending on the same-period chance event of j , and j misreports depending on the same-period chance event of i . However, either of them would not need to be considered if reports were immediate. With this extra requirement, they implemented efficiency in a Perfect Bayesian equilibrium. In contrast, we are using quasi-dominant strategy equilibrium, and we will shortly see why it is an important difference.

We emphasize that the balanced mechanism of Athey and Segal also assume that at the end of each period, it cannot happen that an agent j has already observed some information about the private state of i which i has not observed yet. Also, these periods should be very short in order to utilize all possibilities of making decisions dependent of all earlier chance events of all other players. Therefore, the difference is essentially that Athey and Segal assume this requirement at frequent time points, and we assume it at all time points, hereby avoiding the problems with same-period chance events. This difference does not extend to the (initial) types of the players, just to the changes of the states (chance events) during the game.

In their paper, first they show a mechanism for dependent capability trees but with unbalanced transfers. This is a different direction from ours. Then they consider independent capability trees and balanced transfers, this is more related to our results.

The unbalanced team mechanism is roughly the following. The agents always report all their private information, and according to these reports, the mechanism makes the socially efficient public decision, makes recommendations for the private decisions, and everybody receives the monetary transfer equal to the sum of the valuations of the other agents. They show that truth-telling is a Perfect Bayesian equilibrium.

We show an example that the weakness of this equilibrium concept can easily cause problems. Consider the following setup with two agents. In each round, each agent receives a signal \$1000 or -\$1 independently, with probabilities 1/2. Then they make a public decision “yes” or “no”. If they choose “no”, then both of their valuations are 0. But if they choose “yes”, then their valuations equal the amount in the message. Consider the case when an agent receives signal -\$1. If he reports \$1000 instead, then this costs him at most \$2, but this provides at least \$999 more utility to the other agent. Therefore, reporting \$1000 as long as the other agent also does so is another Perfect Bayesian equilibrium in the infinite-horizon game.

This seems to be not only an abstract counterexample. Similar but less extreme phenomena seems to be very common, unless if the valuations are public information. Namely, whenever the report of an agent has a significant effect on his own reported valuation, then he may report higher valuation in order to build a long-term cooperation relationship with the others, even if his fake report causes some loss in social efficiency.

Their “balanced team mechanism” is different from our “second-price mechanism”, and the equilibrium concepts are also different. Athey and Segal use Perfect Bayesian equilibrium which is a weak concept, and we will show below that in their model, there can also exist inefficient Perfect Bayesian equilibria which seem to be more plausible in practice. In contrast, we use a much stronger and much more convincing equilibrium concept.

To a lesser degree, the same phenomenon may also occur under the balanced team mechanism, which is roughly the following. Whenever an agent i reports a new private signal, the player gets the monetary transfer equal to the expected change on the total expected valuations of all other players during the game. This (signed) monetary transfer is paid by the other agents who share this *equally*.

Consider the following setup with 101 agents including Alice, Bob playing a repeated game. In each round, there is a good to be allocated to one agent. In every round, the true valuation of Alice and Bob of the good are around \$1000, and the valuations of all other agents are clearly lower. In some rounds, Alice knows in advance that her valuation is \$1000; the valuation of Bob is either \$999 or \$1001 with probability 1/2 for each, and he gets to know which one just before the decision. If Bob receives the signal \$1001, but he reports \$999 instead, then Bob will get utility \$500 instead of \$1001 - \$500 = \$501, but Alice gets utility \$1000 - \$500/100 = \$995 instead of \$500/100 = \$5. Assume that Alice and Bob change role alternately. Then reporting \$999 as long as the other agent does so is another Perfect Bayesian equilibrium which is better for Alice and Bob, but not socially efficient.

In contrast, in our mechanism, we do not share the transfer equally, but we do the following. First, we define an arbitrary ordering of the concurring chance events, excluding the reports of the initial states which we keep to be concurrent. The marginal contribution by the reported initial state is compensated by the principal, but for reports of non-initial states of i , each agent $j \neq i$ pays to i the change of j 's own expected total valuation. Hereby, given the capability trees of the agents, we show that the expected utility of each agent with truthful strategy is fixed, even if all other agents collude.

There is one more important difference. They assume publicly known initial states as a parameter of the mechanism. This assumption goes far beyond the common prior assumption, namely, this means that the common prior is assumed to be verifiable by a court. It seems to be very restrictive that the agents receive information dynamically, which the other players cannot completely observe, but the “precise *a priori* distribution” of the entire information flow is verifiable by a court. Especially because the probability distribution of the information flow is a function of the flow of private and public decisions. Moreover, the assumption of independent capability trees means that the common prior must be independent across agents. This makes the assumption even more restrictive. In contrast, in our Project Management Model, all but one players have private capability trees with no prior. (We think of the exceptional player

as the owner of the project.) Our mechanism also satisfies some important requirements: to be individually rational but to avoid free riders. Furthermore, we present a version of our mechanism which is collusion-resistant, but which is less efficient under imperfect competition.

B.2 Comparison to Bergemann–Välimäki

Another related paper was presented by Bergemann and Valimaki (Sep 2006 – 2010) [2], its first version is subsequent to the earlier version of the paper of Athey and Segal (Nov 2004 – 2013) and independent from the earlier version of this paper (Feb 2006 –). Bergemann and Valimaki used an unbalanced dynamic mechanism for social allocations with independent private signals. Their setup is similar to the special case of our setup but without private decisions of the agents, and with discrete time periods. They introduce a mechanism different both from our mechanism and from the mechanism of Athey and Segal, and they show that truthfulness is a Perfect Bayesian equilibrium under that mechanism. This equilibrium has similar weakness as in the result of Athey and Segal, as we show in the example below. In return, their mechanism satisfies an exit condition. This exit condition is a dynamic version of our conditions of individual rationality without offering free lunch. (This dynamic extension is not meaningful in our setting.) The mechanism is roughly the following. In each of an infinite sequence of rounds, the agents are asked to report all of their private information, then we make the socially efficient public decision, and each agent i pays the total expected decrement of the valuations of others caused by considering the preference of i with the decision. Budget is not balanced.

Consider the following setup with at least three agents: Alice, Bob and others. In each round, there are two options: “yes” or “no”, and the rounds are independent. Both Alice and Bob have a preference of about \$1000 for “yes”, meaning that their valuations are \$1000 for “yes” and \$0 for “no”. The other agents prefer “no”, their total preference is about \$2000. The precise amounts will be private information, but their public *a priori* distributions are highly concentrated. Suppose that everybody reports truthfully. If the total preference for “no” is the higher, then this will be the decision, and Alice and Bob will pay nothing. But if “yes” wins by a small amount x , then “yes” will be the decision and both agents pay their reported preference minus x , therefore, both of their utilities will be x . To sum up, the utilities of Alice and Bob are around 0 anyway.

Consider now what happens if Alice, instead of her truthful strategy, reports much more, say, \$3000 for “yes”. If “yes” had won even with true reports, then there is no difference for Alice. If “no” had won by x , then this misreport changed the decision to “yes”, but Alice had to pay her true preference plus x , therefore, she got $-x$ utility instead of 0. This is still a marginal difference. However, in both cases, Bob gets his more favorable decision and he does not have to pay anything, therefore, his utility equals his valuation, which is about \$1000. Therefore, Alice is able to have Bob getting a much higher utility, for at most marginal losses in her own utility. This works vice versa, and if both of them report a much higher preference, then “yes” will be the decision and they will not need to pay anything. Therefore, in the infinite repeated game, it is very likely that Alice and Bob will report higher preferences as long as the other one also does so.

This kind of problem cannot occur in all situations, but still, one should be very careful when applying this mechanism. In contrast, we use a much stronger equilibrium concept, therefore, our mechanism never has such problems.

C Proofs with the offer-based mechanism

Most proofs are either the same or essentially the same under the revelation and the offer-based mechanisms. We show here how to get to the proof of Proposition 9 under the offer-based mechanism. This includes the main differences which appear in the proofs.

For a contract made from a fair-looking proposal, we can define the subcontract from a reported state of the project as the contract corresponding to the subcapability tree (subtree) according to the current state of the project. If all contracts are derived from fair-looking proposals, then we can define the **value $V(T)$ of a reported state of the project T** at a time point after 0 as the principal's maximin utility with the subcontracts from the reported state of the project (and with only these contracts). For an arbitrary $S \subset N^+$, we show that $V(M_S)$ can be calculated here in the same way as with the original definition with the original approach.

The value of the reported starting state (just after 0) is the principal's maximin utility provided that she accepts some agents with the given set of proposals. We can determine the value of all reported states of the project by recursion, using the following.

- The value of a final state is the sum of the valuations at the executions.
- The value of the state of the project changes only at a node; values of states of the project between two neighboring nodes are the same.
- If the values after a decision node are given, then the value at this node is their maximum, because the principal will choose the more favorable option.
- **Lemma 22.** *If the reported state of the project at a chance event is T , the reported states of the project after this chance event are T_1, T_2, \dots, T_k and the reported probabilities are w_1, w_2, \dots, w_k , respectively, then*

$$V(T) = \sum_{i=1}^k w_i \cdot V(T_i).$$

Proof. Let us consider the reported subgame T , and let $x = \sum w_i \cdot V(T_i)$. For calculation, assume that the reported outcome of the chance event will be chosen randomly with the reported probabilities w_i . Furthermore, assume that the principal surely gets the value of the reported state of the project after the chance event; it does not affect the value of any of these reported states of the project. In this case, whichever assignment δ^* she chooses, her expected utility will be $\sum w_i \cdot V(T_i) + \sum w_i \cdot \delta^*(i) = x + 0 = x$, which implies $V(T) \leq x$. On the other hand, $\sum w_i (V(T_i) - x) = 0$, therefore, if the principal assigns $V(T_i) - x$ to the i th branch, then she gets x in all cases. \square

Our recursion also shows that the principal would get the value of the reported starting state as a fixed utility.

D Special cases of the model

D.1 No parallel executions

The messages the principal sends depend only on her not strictly earlier chance events and the previous messages she received. Thus, if an agent $i \in N^+$ is sure that the principal receives no message from anyone else and has no chance event in the time interval $[a, b]$, then, without

decreasing the value of his proposal, i can ask the principal to send him already at a what messages she would send during $[a, b]$ depending on the messages she would have received from i . In addition, if i is sure that the principal will not send any messages to anyone else during $[a, b]$, then, without changing the value of his proposal, i can simplify his offer that he sends all his messages that he would have originally sent during $[a, b]$ only at b .

Consequently, consider a project consisting of two tasks, where the second task can only be started after the first one has been accomplished. We put this into our model in the following way. The consequence provided by each agent working on the first task (called first agent) consists of his completion time C_1 . The consequence provided by of each agent working on the second task consists of his starting time S_2 and the time C_2 he completes; and his capability tree starts with doing nothing until an optional time point S_2 , and then he can start his work. The valuation function of the principal is of the form $\tilde{u}(C_2)$ for a decreasing function $\tilde{u}: \mathbb{R}^{(\text{time})} \rightarrow \mathbb{R}$ if $C_1 \geq S_2$, and $-\infty$ otherwise ($\mathbb{R}^{(\text{time})}$ means the space of time points). The valuation of each agent is simply the minus of his costs. In this case, the principal always communicates only with the agent who is working at the time. Therefore, using the above observation, we can make simplified proposals of the following form, with the same values as of the fair proposals.

If the principal makes \tilde{u} public at the beginning, then the penalty for the chosen second agent is the loss from the delayed completion of the project, therefore, each second agent should demand $\tilde{u}(C_2) - g_2(S_2)$ money, if he can start his work at S_2 and complete it at C_2 . The penalty for the chosen first agent is $-g_2(C_1)$, and each first agent declares how much money he demands for the first task depending on the penalty function. Then the principal chooses the pair by which she gains the highest utility.

Formally, the form of the simplified proposal for the first agents is a function $g_1: (\mathbb{R}^{(\text{time})} \rightarrow \mathbb{R}) \rightarrow \mathbb{R}$, and for the second agents this is a function $g_2: \mathbb{R}^{(\text{time})} \rightarrow \mathbb{R}$. If all proposals are so, then the principal chooses a pair for which $g_1(g_2)$ is the greatest. Then she tells the penalty function g_2 to the chosen first agent at the beginning of his capability tree, and, after his completion, she pays him $g_2(C_1) - g_1(g_2)$. Then the chosen second agent can start his work at C_1 and after his completion, he gets $\tilde{u}(C_2) - g_2(C_1)$. This way, the principal gets utility $\tilde{u}(C_2) - (g_2(C_1) - g_1(g_2)) - (\tilde{u}(C_2) - g_2(C_1)) = g_1(g_2)$.

In the simplified fair proposals, g_1 and g_2 are chosen in such a way that makes their expected utility independent of the arguments (h and C_2 , resp.), if the agents use their best strategies afterwards.

If a first agent has no choice in his capability tree, that is, his completion time C_1 is simply a probabilistic variable, then he should choose $g_1(h) = \mathbf{E}(h(C_1)) - x$, where x is his costs plus his profit.

D.2 Controlling and controlled players

For an example, consider a task of building a unit of railroad. An agent i can make this task for a cost of 100, but with 1% probability of failure which would cause a huge loss of 10,000. Another agent j could inspect and, in the case of failure, correct the work of i under the following conditions. The inspection costs 1. If the task was made correctly, then he does nothing else. If not, then he detects and corrects the failure with probability 99% for a further cost of 100. But with probability 1% he does not detect the failure, and therefore does nothing. If both of them use truthful strategy and they are the accepted agents for the task, then the mechanism works in the following way.

In the end, i gets 101.99, but pays 199 (totally he pays 97.01) compensation if he fails. j gets 1 if he correctly finds the task to be correct, and, beyond this, he gets 200 if the task was made badly but he corrects it, but he pays 9800 if he misses correcting it.

It can be checked that the expected utility of each agent is his profit independently of the

behavior of the others, and that the utility of the principal is fixed.

E Observations for application

From now on, we use the offer-based mechanism.

E.1 Modifications during the process

In practice, the capability trees can be extremely difficult, therefore, submitting precise fair proposals cannot be expected. Hence players can only present a simplified approximation of their capability tree. Generally, such inaccuracies do not significantly worsen the optimality; nevertheless, this loss can be reduced far more with the following observation.

Assume that someone, whose proposal has been accepted, can refine his capability tree during the process. It would be beneficial to allow him to modify his proposal correspondingly. The question is: on what conditions?

The answer for us is to allow him to modify his capability tree if he pays the difference between the maximin utility of the principal with the original and the new proposals. Or, equivalently, the contract as a payment function automatically decreases by this difference. It is easy to see that, for an agent, whether and how to modify his proposal is the same question as which proposal to submit among the proposals with a given value. Consequently, Theorem 19 shows that the agent has incentives to change to his true fair proposal.

With the original mechanism, this is equivalent to the following. Each possible modification can be handled as a chance node in the original proposal with 1 and 0 probabilities for continuing with the original and the modified proposal, respectively. Because at such a chance node, the principal assigns 0 to the branch of not modifying; and to the modification, she assigns the difference between the values of the reported states of the project after and before.

It may happen that at the beginning, it is too costly for some agent to explore the many improbable branches of his decision tree, especially if he does not yet know whether his proposal will be accepted; but later, it would be worth exploring better the ones that became probable. These kinds of in-process modifications are what we would like to make possible. We show that each player has approximately the same incentives as the “total incentives” of the players in the better scheduling of these small modifications.

The expected utility of an agent with an accepted fair proposal is fixed and for a nearly fair agent, the little modifications of the other proposals have negligible effect. As the modifications of each agent have no influence on the utility of the principal and only this negligible influence on the expected utility of other agents, the change in the expected total utility is roughly the same as the change in the expected utility of this agent. This confirms the above statement.

A very similar argument shows that, under the offer-based mechanism, the principal can also do such in-process modifications, and her incentives in doing this is about as much as the total incentive of all players.

E.2 Risk-averse agents

Assume that an agent i has a non-quasi-linear utility function $u_i(x_i, t_i)$, where x_i is the execution of the capability tree of i . We assume that $\forall x: \lim_{t \rightarrow \infty} u_i(x, t) = \infty$. For example, if $u_i = w_i(v_i + t_i)$ with a monotone increasing concave function w_i , then this expresses risk aversion. We define a proposal to be **reasonable** in the same way as the fair proposal with the only difference being

that the agent demands a transfer t_i satisfying

$$u_i(\hat{\xi}_i, t_i) \geq \sum_{\chi \in X_i} \delta^*(\chi)$$

By a reasonable proposal, in the case of acceptance, the expected utility of the valuation of the agent is independent of the choices of the principal. If all proposals are reasonable, then the utility of the principal remains fixed. If the agent is risk-neutral, then his reasonable proposal is fair. These are some reasons why reasonable proposals work “quite well”. We do not state that it is optimal in any sense, but a reasonable proposal may be better than a fair proposal in the risk-averse case.

We note that the evaluation of reasonable proposals can be much more difficult than of fair proposals, but for each agent i , if $h_i(x) = a_i - b_i \cdot e^{-\lambda_i x}$, then a similar recursion works as in Lemmas 6 and 7.

E.3 Necessity of being informed about one’s own process

We assumed that none of the chosen players knew anything better about a chance event of any other chosen agent. We show an example that does not satisfy this assumption and the mechanism is inefficient. Consider two agents i and j that will definitely be accepted. Assume that i believes the probability of an unfavorable event in his work to be 50%, but j knows that the probability is 60%, he knows the estimation of i , and he also knows that at a particular decision node of his, he will be asked to make a move corresponding to the outcome of this chance event of i . It can be checked that if the proposal of i is fair, then if j increases the demanded payment in his proposal of the more probable case by an amount of money and decreases it in the other case by the same amount, then the value of his proposal remains the same, but this way, j bets 1 : 1 with i on an event of 60% probability.

However, if the assumption *almost* holds, then we can apply some techniques which limits these problems.

If an agent i has only a small risk that his information might not dominate the belief of someone else about i ’s probabilities, then in order to limit potential losses, he could rightfully say that a larger bet can only be increased on worse conditions. Submitting a reasonable proposal with risk-averse utility function makes something similar, which is another reason to use this.

Another technique is that the agent reports slightly less flexibility. In the example in the first paragraph, this means that after i gets to know that one of the events would provide him larger payment, he uses his unreported capability to increase the probability of this event.

All these problems can be completely avoided when the chance event is contractible. In this case, we can even “backdate” the chance event as long as the player can prove that he did not have influence on its probability between its earlier and actual dates. In this case, transfers are calculated so that we keep the stochastic strategy of the reported game according to the actual date of the reports of the chance events, but we consider the changes in the expected total utility by the reported chance events according to the new chronological ordering.

Finally, if spying about the working processes of others are possible but difficult (e.g. costly), then many different reputation systems could be used to exclude permanent cheaters, based on correlation analysis in previous projects. For the simplest example, we could allow proposals to be conditioned on which other agents will be accepted.

E.4 Agents with limited liability

This subsection is only a suggestion for the cases when we have some agents with limited liability, and we cannot provide any clear mathematical statement about its level of efficiency.

Our mechanism requires each agent to be able to pay as much money as the maximum possible damage he could have caused. However, in many cases, there may be some agents who cannot satisfy this requirement. Despite this, accepting such an agent i may still be a good decision, if i is reliable to some degree.

To solve this problem, i should find someone who has enough funds, and who takes the responsibility, for example for an appropriate fee. If the agent is reliable to some degree, then he should be able to find such insurer player j . (It can even be the principal, but is considered as another player.) This method may also be useful when i has enough funds, but is very risk-averse.

Here, i and j work similarly as the controlled and controlling parties in Appendix D.2. The differences are that j does not work here, but he has some knowledge about the capability trees (mainly about the capability tree of i). Furthermore, the role of j here can be combined with his role in Appendix D.2.

E.5 Simplified capability trees

In many situations, this mechanism cannot be applied directly because of its complex administrative requirements. For example, consider a supermarket employing many cashiers. Some of the cashiers can go to work whenever they are asked to, while others need to get to know their schedule well in advance. And everyone can be ill, etc. Clearly, it would be unrealistic to ask them to precisely define their future as a capability tree.

However, it might be useful to construct a weaker but simpler class of proposals which require very few and simple communications, but which could still express the preferences and requests quite well. Such a mechanism might still be much more efficient than the naive mechanisms.

We show an example about how the author believes this would work in practice. A cashier who reports a level of permanent availability would get a base salary X , but he may receive messages like “We need you tomorrow from 7:00 to 15:00 for an increased salary Y , but you get Z deduction if you do not come.” X , Y and Z are chosen in a fair way, depending on the level of availability the cashier reported. And of course, everyone can report modifications in their availability (e.g. via an online system) with the fair conditions defined by the mechanism.

Another example is about contracting for gas or electricity. In practice, some companies get cheaper electricity if they accept that in the case of problems in the current supply, they may be switched off. But if the companies gave a rough description of their incentives about it, and the electric supplier would also make a dynamic stochastic description of the possible problems, then we would be able to make more efficient decisions for the economy.

F A general framework

We present a framework for the analysis of first and second-price and other mechanisms on a larger class of problems. Therefore, we redefine the notions and notations, but these remain analogous to which we used so far.

When we apply it to the Project Management Model, the action of a player should be interpreted as his strategy given his own type. We may exclude some strategies from the “action sets” which are not rational *after* the first move.

Let $N = \{0, 1, 2, \dots, n\}$, where 0 represents the principal and $N^+ = \{1, 2, \dots, n\}$ is the set of agents. We have a prior distribution μ on Θ^N , a set A of feasible actions/offers, $\emptyset \in A$, and two functions (valuation) $v: A^N \rightarrow \mathbb{R}$ and (utility) $u: \Theta^N \times A^N \rightarrow \mathbb{R}^N$. We assume that

$$\forall \theta \in \Theta^N, \forall i \in N^+, \forall \mathbf{a}_{-i} \in A^{N \setminus \{i\}}: u_i(\theta, (\mathbf{a}_{-i}, \emptyset)) = 0,$$

which will express that rejected players get utility 0. There is a function (decision) $d : A^N \rightarrow A^N$ satisfying that

$$\forall \mathbf{a} \in A^N: \quad d(\mathbf{a}) \in \arg \max_{\mathbf{b} \in \{a_0\} \times \{\emptyset, a_1\} \times \{\emptyset, a_2\} \times \dots \times \{\emptyset, a_n\}} v(\mathbf{b}).$$

The strategy set of player i is the set of functions $S_i = \{\Theta \rightarrow A\}$. Assume that for each $i \in N$, there exists a (truthful) strategy $\tau_i \in S_i$ satisfying the following conditions.

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N^+, \forall \mathbf{a}_{-i} \in A^{N \setminus \{i\}}: \quad u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, \tau_i(\theta_i))) = 0$$

$$\forall \mathbf{a}_{N^+} \in A^{N^+}: \quad u_0(\boldsymbol{\theta}, (\mathbf{a}_{N^+}, \tau_0(\theta_0))) = v(d(\mathbf{a}_{N^+}, \tau_0(\theta_0))).$$

$$\forall \boldsymbol{\theta} \in \Theta^N, \forall \mathbf{a} \in A^N: \quad \sum_{i \in N} u_i(\boldsymbol{\theta}, d(\mathbf{a})) \leq v(d(\boldsymbol{\tau}(\boldsymbol{\theta})))$$

Assume that for each $r \in A$ and $x \in \mathbb{R}$, there exists an offer denoted by $r + x \in A$ satisfying $\forall \boldsymbol{\theta} \in \Theta^N, \forall i \in N^+, \forall \mathbf{a}_{-i} \in A^{N \setminus \{i\}}:$

$$v(\mathbf{a}_{-i}, r + x) = v(\mathbf{a}_{-i}, r) - x,$$

$$u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, r + x)) = u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, r)) + x.$$

These were the main assumptions. Optionally we might have further assumptions, we will come back to them, but let us see now how to use the framework.

Let us define the marginal contribution of an offer a_i of a player $i \in N$ by

$$v_i^+(\theta_0, \mathbf{a}) = v(\theta_0, d(\theta_0, \mathbf{a})) - v(\theta_0, d(\theta_0, \mathbf{a}_{-i}, \emptyset)).$$

A strategy profile $\mathbf{s} = (s_1, s_2, \dots, s_n) \in \mathbf{S} = S_1 \times S_2 \times \dots \times S_n$ is an equilibrium under the first-price mechanism if

$$\forall i \in N^+, \forall s'_i \in S_i: \quad \mathbb{E}_\mu \left(u_i(\boldsymbol{\theta}, d(\mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}), s'_i(\theta_i))) \right) \leq \mathbb{E}_\mu \left(u_i(\boldsymbol{\theta}, d(\mathbf{s}(\boldsymbol{\theta}))) \right).$$

A strategy profile $\mathbf{s} \in \mathbf{S}$ is an equilibrium under the second-price mechanism if

$$\forall i \in N^+, \forall s'_i \in S_i:$$

$$\mathbb{E}_\mu \left(u_i(\boldsymbol{\theta}, d(\mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}), s'_i(\theta_i))) + v_i^+(\mathbf{s}_{-i}(\boldsymbol{\theta}_{-i}), s'_i(\theta_i)) \right) \leq \mathbb{E}_\mu \left(u_i(\boldsymbol{\theta}, d(\mathbf{s}(\boldsymbol{\theta}))) + v_i^+(\mathbf{s}(\boldsymbol{\theta})) \right),$$

$$\begin{aligned} \text{and } \forall s'_0 \in S_0: \quad & \mathbb{E}_\mu \left(u_0(\boldsymbol{\theta}, d(\mathbf{s}_{N^+}(\boldsymbol{\theta}_{N^+}), s'_0(\theta_0))) - \sum_{i \in N^+} v_i^+(\mathbf{s}_{N^+}(\boldsymbol{\theta}_{N^+}), s'_0(\theta_0)) \right) \\ & \leq \mathbb{E}_\mu \left(u_0(\boldsymbol{\theta}, d(\mathbf{s}(\boldsymbol{\theta}))) - \sum_{i \in N^+} v_i^+(\mathbf{s}(\boldsymbol{\theta})) \right). \end{aligned}$$

We can generalize Theorems 4 and 21 under this general framework, because the proofs used only properties and assumptions what we stated above. Therefore, this framework can be useful for further analysis about the level of efficiency of the first-price mechanism. Furthermore, we can define the generalized version of the conjecture mentioned in Section 6.1.

Conjecture 23. For each common prior μ , there exists an equilibrium $\mathbf{s} \in \mathbf{S}$ under the second-price mechanism satisfying

$$\mathbb{E}_\mu \left(u_0(\boldsymbol{\theta}, d(\boldsymbol{\tau}(\boldsymbol{\theta}))) - \sum_{i \in N^+} v_i^+(\boldsymbol{\tau}(\boldsymbol{\theta})) \right) \leq \mathbb{E}_\mu \left(u_0(\boldsymbol{\theta}, d(\mathbf{s}(\boldsymbol{\theta}))) - \sum_{i \in N^+} v_i^+(\mathbf{s}(\boldsymbol{\theta})) \right).$$

Now we define further potential assumptions which might be useful and which were satisfied in our model.

We may assume that whether a player wins with an offer or with a constant higher offer has no effect on the other players:

$$\forall j \in N \setminus \{i\}: \quad u_j(\boldsymbol{\theta}, (\mathbf{a}_{-i}, a_i + x)) = u_j(\boldsymbol{\theta}, \mathbf{a}).$$

We may assume that the set of offers is convex, namely, $\forall p, q \in A$ and $\lambda \in (0, 1)$, there exists a strategy denoted by $\lambda p + (1 - \lambda)q \in A$ satisfying that $\forall \boldsymbol{\theta} \in \Theta^N$, $\forall i \in N$, $\forall \mathbf{a}_{-i} \in A^{N \setminus \{i\}}$:

$$v(\theta_0, (\mathbf{a}_{-i}, \lambda p + (1 - \lambda)q)) = \lambda \cdot v(\theta_0, (\mathbf{a}_{-i}, p)) + (1 - \lambda) \cdot v(\theta_0, (\mathbf{a}_{-i}, q)),$$

$$u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, \lambda p + (1 - \lambda)q)) = \lambda \cdot u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, p)) + (1 - \lambda) \cdot u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, q)).$$

We may assume that each player can submit multiple offers, namely, $\forall p, q \in A$, there exists an offer denoted by $p \wedge q \in A$ satisfying that $\forall \boldsymbol{\theta} \in \Theta^N$, $\forall i \in N$, $\forall \mathbf{a}_{-i} \in A^{N \setminus \{i\}}$:

$$v(\theta_0, (\mathbf{a}_{-i}, p \wedge q)) = \max_{r \in \{p, q\}} \left(v(\theta_0, (\mathbf{a}_{-i}, r)) \right),$$

and with the very same r ,

$$u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, p \wedge q)) = u_i(\boldsymbol{\theta}, (\mathbf{a}_{-i}, r)).$$

Furthermore, we may assume that v is independent from a_0 . Or we may assume that the principal is forced to choose a_0 from a restricted set $A_0(\theta_0)$, and v is independent of this restricted choice of a_0 . Or we may assume that $|A_0| = 1$ meaning that the principal has no strategic capacity.

Index

- $M = \Gamma/m_1$: project Γ under the first-price mechanism m_1 , 19
- $V(G)$: cooperative value of a game G , 16
- $V(\boldsymbol{\theta}) = V(\Gamma(\boldsymbol{\theta}))$: joint value of a vector of capability trees $\boldsymbol{\theta}$, 17
- $[\emptyset]$: empty capability tree, 14
- W : set of accepted agents, 34
- $W_0 = W \cup \{0\}$: set of accepted agents and the principal, 36
- X_i : set of all chance events during the execution of i 's capability tree, 17
- Γ : project, 14
- $\Gamma(\boldsymbol{\theta})$: Γ with capability trees $\boldsymbol{\theta}$, 15
- $\Gamma(\hat{\boldsymbol{\theta}})$: reported project, 17
- $N = \{0, 1, 2, \dots, n\}$: set of all players, 14
- $N^+ = \{1, 2, \dots, n\}$: set of all agents, 14
- \mathcal{S}_i : strategy set of player i , 11, 14
- χ : chance event, 17
- τ_i : truthful strategy of player i , 19
- $\delta(\chi)$: change in the expected total utility in the reported project, 17
- φ : strategy form, 35
- ξ_i : execution of player i , 11, 13
- $\hat{\xi}_i$: reported execution of player i , 17
- θ_i : capability tree of player i , 2, 13
- $\hat{\theta}_i$: proposal (reported capability tree), 17
- p_i : payment to player i , 14
 - first-price payment, 17
- $p_i^{2nd} = p_i + v_i^+(\hat{\boldsymbol{\theta}})$: second-price payment, 18
- s_i : strategy of player i , 11
- \hat{s} : an efficient strategy profile of the reported project $\hat{\boldsymbol{\theta}}$, 17
- u_i : utility of player i , 15
- u_S : total utility of all players in $S \subseteq N$, 16
- u_i^{2nd} : utility of player i under the second-price mechanism, 19
- $v_i^+ = v_i^+(\hat{\boldsymbol{\theta}})$: value of the proposal of i , 17
- v_i^\pm : signed value of a proposal of i , 35
- \hat{v}_i : reported valuation function of i , 17
- w_i : reported probability (or weight), 24
 - action, 14
 - agent $i \in N^+$
 - accepted agent $i \in W$, 34
 - rejected agent $i \in N^+ \setminus W$, 34
 - Bayesian belief system, 31
 - capability tree θ_i , 2, 13
 - empty capability tree $[\emptyset]$, 14
 - omissible capability tree, 13
 - chance event, 13
 - chance node, 13
 - collusion-resistance, 16
 - cooperative value V , 16
 - execution of a capability tree ξ_i , 11, 13
 - reported execution $\hat{\xi}_i$, 17
 - free lunch, 16
 - game form, game, 11
 - subgame, 12
 - mechanism, 16
 - first-price mechanism, 4, 16
 - offer-based mechanism, 5, 18
 - revelation mechanism, 4, 16
 - second-price mechanism, 4, 17
 - payment p_i , 14
 - perfect competition, 22, 33
 - planner, 14
 - principal (player 0), 14
 - project Γ , 14
 - reported project $\Gamma(\hat{\boldsymbol{\theta}})$, 17
 - proposal (reported capability tree) $\hat{\theta}_i$, 4, 17
 - fair proposal $\theta_i - x$, 4, 23
 - quasi-dominant equilibrium, 20, 29
 - weak quasi-dominant eq., 30, 39
 - signed value of a proposal v_i^\pm , 35
 - state: identified with subgame, 3, 12
 - strategy s_i
 - fair strategy $\tau_i - x$, 23
 - fair-looking strategy, 35
 - truthful strategy τ_i , 19
 - strategy form φ , 35
 - fair strategy form $\varphi(\tau_i)$, 35
 - total utility u_N , 16
 - utility u_i , 15
 - maximin utility, 18
 - valuation $v_i = v(\xi_i)$, 13
 - valuation function v_i
 - reported valuation function \hat{v}_i , 17
 - value of a proposal v_i^+ , 17
 - value of a state, 4, 46
 - value of the strategy form $V_i(\varphi(s_i))$, 35