Individually Rational Strategy-Proof Social Choice with Exogenous Indifference Sets

Mingyu Guo[1], Yuko Sakurai[2(✉)], Taiki Todo[2], and Makoto Yokoo[2]

[1] The University of Adelaide, Adelaide, Australia
mingyu.guo@adelaide.edu.au
[2] Kyushu University, Fukuoka, Japan
{ysakurai,todo,yokoo}@inf.kyushu-u.ac.jp

Abstract. We consider a social choice problem where individual rationality is required. The status quo belongs to the outcome space, and the selected alternative must be weakly better than the status quo for everybody. If the mechanism designer has no knowledge of the alternatives, we obtain a negative result: any individually rational (IR) and strategy-proof (SP) mechanism can choose at most one alternative (besides the status quo), regardless of the preferences. To overcome this negative result, we consider a domain where the alternatives have a known structure, i.e., an agent is indifferent between the status quo and a subset of the outcomes. This set is exogenously given and public information. This assumption is natural if the social choice involves the participation of agents. For example, consider a group of people organizing a trip where participation is voluntary. We can assume each agent is indifferent between the trip plans in which she does not participate and the status quo (i.e., no trip). In this setting, we obtain more positive results: we develop a class of mechanisms called Approve and Choose mechanisms, which are IR and SP, and can choose multiple alternatives as well as the status quo.

1 Introduction

Social choice theory, which is a primary research field in micro-economics, studies the design and analysis of mechanisms/rules for collective decision making. Recently, due to the growing needs for agent technology in terms of complex information systems, various studies on social choice theory have been conducted in artificial intelligence and multi-agent systems [1,4,10].

We consider *individually rational* (IR) and *strategy-proof* (SP) mechanisms for social choice settings where the *status quo* belongs to the outcome space. Agents submit their preference orders over all outcomes including the status quo. By IR, the selected outcome must be weakly better than the status quo for everybody. IR also implies that the agents have veto power over the alternatives; if an agent declares that one alternative is worse than the status quo, it cannot be chosen. Analyzing veto power for a social choice function is an important problem in social choice literature [3,12].

M. Baldoni et al. (Eds.): PRIMA 2016, LNAI 9862, pp. 181–196, 2016.
DOI: 10.1007/978-3-319-44832-9_11

Unfortunately, we derive a negative result for general settings. If a mechanism is IR and SP, it must choose one particular alternative x in advance and choose between the status quo and x based on the agent preferences. Even if there exist many alternatives as well as the status quo, all but one will never be chosen, regardless of the agents' preferences. This is highly restrictive.

In light of the above negative result, we introduce a domain where the alternatives have a known structure. This allows us to design IR and SP mechanisms that can determine an outcome among multiple alternatives besides the status quo. We assume that for each agent, a set of alternatives, which are equivalent to the status quo for her (called the *indifference set*), is exogenously given. This information is public. For each agent, we call a choice that is not in her indifference set a standard alternative. An alternative is called a *multiagent choice* if multiple agents regard it as a standard alternative. We show that when the number of agents is n, any IR and SP mechanism can determine an outcome among at most $n-1$ multiagent choices. Furthermore, we develop a new class of IR and SP mechanisms called *Approve and Choose* (AC) mechanisms that can choose up to $n-1$ multiagent choices.

When there exist only two agents (agent 1 and 2), the AC mechanism first chooses a set of alternatives $X(f)$ in some way independently from their declarations. We assume $X(f)$ includes the status quo and exactly one multiagent choice x^*; other alternatives $X(f)\backslash x^*$ are standard alternatives only to agent 1. The mechanism first ask agent 2 whether x^* is acceptable, i.e., x^* is better than the status quo. If x^* is acceptable, then agent 1 selects her most preferred alternative from $X(f)$. Otherwise, she selects her most preferred alternative from $X(f)\backslash\{x^*\}$. It is clear that this mechanism is SP and IR. This idea can be generalized to n-agent cases.

We also show the modified AC mechanisms so that the mechanisms can be applied to settings with *quasi-linear utilities*. For such settings, we show that there exist AC mechanisms that simultaneously guarantee individual rationality, strategy-proofness, and *strong budget-balancedness*.

These assumptions we apply in this paper are appropriate in many application domains. For example, assume a travel agency is arranging a group trip. Each alternative can be associated with a different venue and a different set of participants. We can assume that each person is indifferent among tours in which she does not participate. The travel agency can choose one alternative using the AC mechanism. Similarly, consider a programming contest organized by crowdsourcing. Each alternative (a project's candidate) can be associated to a different goal and a different set of programmers. It is reasonable to assume that a programmer is indifferent about the projects she is not involved in. Furthermore, when deciding the location of a facility in one of two relatively remote cities R_1 and R_2, R_1's residents care about the exact location of the facility when it is located in R_1, but they can be indifferent about the exact location when it is located in R_2 since it is already far away.

2 Related Works

We introduce several recent works in social choice theory related to this paper.

Faltings [6] proposed SP and budget-balanced mechanisms that work by sacrificing efficiency for quasi-linear utility settings. Lu and Boutilier considered the multi-winner social choice problem when voter preference profiles are incomplete using the notation of the minimax regret and proposed a greedy algorithm to determine a robust slate of options. However, these models do not assume the existence of the status quo. If we apply these mechanisms in our setting, we cannot guarantee IR.

Sato [13] used an idea called exogenous indifference classes that closely resembles our indifference set. Such indifference classes of agent preferences are exogenously given. He focused on set A of alternatives that every agent strictly ranks, and showed that strategy-proofness implies a dictatorship when $|A| \geq 3$. Since our newly developed mechanism deliberately chooses alternatives utilizing public information so that $|A| \leq 2$ holds, it is SP and non-dictatorial.

Barberá and Ehlers [2] also considered indifferences in agent preferences in voting situations and clarified the necessary and sufficient condition for the majority rule to be quasi-transitive. In this sense, their analysis is quite different from ours; we focused on strategy-proofness.

Darmann *et al.* [5] addressed a group activity selection problem among agents, where an agent's utility of an activity can depend on the number of participants. Our model is more general; the utility can depend on other factors, e.g., who are going to participate.

Guo *et al.* [9] proposed the optimal shifted Groves mechanism, which is IR, SP, and non-deficit, and minimizes the worst-case efficiency loss when choosing an outcome from a finite set including the status quo in a quasi-linear domain. To be precise, they require that there exists at least one alternative for which the total valuation of the agents is nonnegative, which is more lenient than requiring the status quo. They assume that every agent's valuation for every outcome is bounded and the bounds are public information. We use a different model in which the set of alternatives over which each agent is indifferent is public information.

3 Preliminaries

A social choice problem is defined by tuple $(N, X, q, \succ)$. $N = \{1, 2, \dots, n\}$ is a set of agents, and X is a finite set of alternatives. We assume $n \geq 2$. Otherwise, choosing one alternative is easy; the agent can only choose her favorite alternative. Alternative $q \in X$ is called the status quo. $\succ = (\succ_i)_{i \in N} \in \mathcal{P}^n$ is a profile of agent preferences over X. Here, $\mathcal{P}$ is the set of all possible preference relations for an agent. For any pair of alternatives, $x, x' \in X$, $x \succ_i x'$ means agent i strictly prefers x over x'. We assume each $\succ_i$ is transitive, complete, and anti-symmetric. Let N_{-i} denote $N \setminus \{i\}$.

A social choice function (or a mechanism) selects one of the alternatives given in the preference profile of the agents.

Definition 1 (Social Choice Function). *A social choice function is a function* $f : \mathcal{P}^n \to X$*.*

Let $X(f)$ denote $\{x \in X : \exists \succ \in \mathcal{P}^n \text{ s.t. } f(\succ) = x\}$, where $X(f) \subseteq X$ is the set of all the alternatives, each of which has a chance to be selected.

Next, we introduce several desirable properties of a social choice function.

Definition 2 (Individual Rationality). *We say social choice function* f *is individually rational (IR) if for all* $i \in N$ *and for all* $\succ \in \mathcal{P}^n$, $f(\succ) \succ_i q$ *or* $f(\succ) = q$ *holds.*

Individual rationality means that the selected alternative must be at least as good as the status quo. Note that we do not assume that the status quo is the worst alternative for each agent; an agent may prefer the status quo over alternative x.

Definition 3 (Strategy-Proofness). *We say social choice function* f *is strategy-proof (SP), if for all* $i \in N$*, for all pairs of* $\succ, \succ' \in \mathcal{P}^n$ *such that* $\succ_j = \succ'_j$ *for any* $j \in N_{-i}$, $f(\succ) \succ_i f(\succ')$ *or* $f(\succ) = f(\succ')$ *holds.*

We assume $\succ_i$ is the private information of agent i. Thus, the outcome is calculated based on declared preferences. If f is not SP, an agent has an incentive to misreport her preference.

4 Impossibility for General Cases

We show a negative result for general cases. The only mechanism that satisfies desirable properties can choose at most one alternative besides the status quo.

Theorem 1. *If* f *is IR and SP, then* $|X(f)| \leq 2$ *holds, i.e.,* f *can choose at most two alternatives.*

Proof. From the Gibbard-Satterthwaite theorem [7,14], any SP mechanism must be dictatorial if the number of alternatives is more than or equal to 3. However, a dictatorial mechanism does not satisfy IR. Thus, $|X(f)|$ is at most 2.

It is clear that $q \in X(f)$ must hold as long as f satisfies IR. If a mechanism is IR and SP, then it is either trivial (always chooses q) or it chooses between q and one particular alternative x, which is selected in some way (perhaps in arbitrarily) independently from the agent declarations. The mechanism chooses x if all agents prefer it over q. Otherwise, it chooses q.

5 Social Choice with Exogenous Indifference Sets (SC-EI)

To overcome the negative result obtained in the previous section, we consider a domain where the alternatives have a publicly known structure: a *Social Choice with Exogenous Indifference sets* problem (SC-EI).

An SC-EI is defined by a tuple $(N, X, q, \succ, Q)$. The meanings of $N, X, q, \succ$ are identical to the social choice problem defined in the previous section. Here $Q = (Q_i)_{i \in N}$ is a profile of the indifference sets. Each $Q_i \subseteq X$ represents the set of alternatives that are equivalent to q for agent i where for any $x \in Q_i$, $x \sim_i q$ holds. We assume $\sim_i$ is reflective, transitive and symmetric. We assume $q \in Q_i$ holds for each $i \in N$. We also assume the preference of each agent is strict except for the elements of Q_i, i.e., for any $x \in X \setminus Q_i$ and $x' \in X$, where $x \neq x'$, either $x \succ_i x'$ or $x' \succ_i x$ holds.[1] We denote $x \succeq_i x'$ when $x \succ_i x'$ or $x \sim_i x'$ holds. We say x is indifferent to agent i (with q) if $x \in Q_i$. x is also a standard alternative for agent i if $x \in X \setminus Q_i$.

We assume Q is public information and develop a mechanism that exploits this information. Assuming Q is public is reasonable if each alternative is relevant to only a subset of agents, e.g., an agent is indifferent between alternative trip plans in which she does not participate.

Let us show a concrete example and an informal description about how a mechanism can utilize Q.

Example 1. Consider the following case.

- $N = \{1, 2, 3\}$, i.e., three agents.
- $X = \{x_1, x_2, x_3, x_4, x_5, q\}$, i.e., six alternatives including the status quo. Here, we assume these alternatives are possible trip plans, where x_1 is a plan that all agents go to Thailand, x_2 is a plan that all agents go to Singapore, x_3 is a plan that only agents 1 and 2 go to Thailand, x_4 is a plan that only agents 1 and 3 go to Singapore, and x_5 is a plan that only agent 1 goes to Thailand.
- $Q_1 = \{q\}, Q_2 = \{x_4, x_5, q\}, Q_3 = \{x_3, x_5, q\}$, i.e., each agent considers any trip that she does not participate in as an element of her indifference set.

If Q is not public knowledge, as shown in the previous section, when $|X(f)| = 2$, mechanism f first chooses one alternative x within $\{x_1, \ldots, x_5\}$ in some way, e.g., arbitrary. Then, the mechanism asks each agent whether she approves it, i.e., whether she prefers x over q. If all the agents approve x, then it is chosen. Otherwise, q is chosen.

When Q is public, instead of choosing one alternative, the mechanism can choose two, say, x_3 and x_4. Here, since $x_3 \in Q_3$, it can select x_3 without the approval of agent 3; it is public knowledge that she is indifferent between x_3 and q. The mechanism asks agent 2 whether she approves x_3. Also, it asks agent 3 whether she approves x_4. Let $\widehat{X}$ denote the set of approved alternatives. Then, the mechanism lets agent 1 choose her most preferred alternative within $\widehat{X} \cup \{x_5, q\}$. Clearly, $X(f) = \{x_3, x_4, x_5, q\}$. Thus, $|X(f)|$ becomes strictly more than 2.

We next introduce several criteria to evaluate social choice functions. For alternative $x \in X$, let $S(x)$ denote $\{i \in N : x \notin Q_i\}$, i.e., the set of all agents

[1] This assumption is required to obtain Maskin monotonicity [11], which is a powerful tool to show various properties of a mechanism.

who consider x a standard alternative. In Example 1, $S(x_1) = S(x_2) = \{1, 2, 3\}$, $S(x_3) = \{1, 2\}$, $S(x_4) = \{1, 3\}$, and $S(x_5) = \{1\}$.

Definition 4 (Multiagent Choice). *We say alternative x is a multiagent choice if $|S(x)| \geq 2$. Also, we say x is a multiagent choice of mechanism f if x is a multiagent choice and $x \in X(f)$ holds. Let $M \subseteq X$ denote the set of all multiagent choices in X and let $M(f)$ denote $X(f) \cap M$, i.e., the set of all multiagent choices with a chance to be selected by mechanism f. We call $M(f)$ the possible multiagent choices of mechanism f.*

In Example 1, x_1, x_2, x_3, and x_4 are multiagent choices. Also, for the mechanism described in Example 1, $M(f) = \{x_3, x_4\}$. Let $X_i(f)$ denote $X(f) \backslash Q_i$ and let $M_i(f)$ denote $M(f) \cap X_i(f)$.

In the trip organizing example, a multiagent choice is a trip involving more than one agent. It is natural to assume that a social choice function is better if it can choose more multiagent choices. In the general setting, we showed that at most one alternative besides the status quo can be chosen as long as mechanism f is IR and SP. Thus, the number of possible multiagent choices of f is at most 1. In an SC-EI, the number of possible multiagent choices can be increased, e.g., in Example 1, $|M(f)| = 2$. In the next section, we show that the number of possible multiagent choices is bounded, i.e., at most $n - 1$.

From the above definition, for x to be a multiagent choice, $|S(x)| \geq 2$ is sufficient; it does not care whether $|S(x)| = 2$ or $|S(x)| = n$. In the trip organizing example, it is reasonable to assume that a trip plan with more participants is desirable. Thus, let us introduce another criterion that takes into account the quantity of agents who consider an alternative to be standard.

Definition 5 (Agent Count). *For subset of multiagent choices $M' \subseteq M$, we call $C(M') = \sum_{x \in M'} |S(x)|$ its agent count. We also call $C(M(f))$ the agent count of mechanism f.*

In Example 1, $C(M(f)) = |\{1, 2\}| + |\{1, 3\}| = 4$. This value means the cumulative total number of participating agents of trips that involve multiple agents.

In general settings, the agent count is at most n, since there exists at most one multiagent choice. In an SC-EI, the number of multiagent count can be increased, e.g., in Example 1, $C(M(f)) = 4$. In the next section, we show that $C(M(f))$ is bounded, i.e., at most $2(n - 1)$.

6 Properties of SC-EI

We show the bounds of the size of the possible multiagent choices and the agent count of an IR and SP mechanism. Let us introduce another assumption called weak non-bossiness. Let $L(x, \succ_i) = \{y \in X : x \succeq_i y\}$. $L(x, \succ_i)$ is called agent i's *weak lower contour set* for alternative x. In words, the lower contour set of x of agent i means a set of alternatives that are less preferred than or equal to x for i.

Definition 6 (Weak Non-bossiness). *We say social choice function f is weakly non-bossy (WNB) if for any $i \in N$ and any $\succ, \succ' \in \mathcal{P}^n$ s.t. for all $j \in N_{-i}$, $\succ_j = \succ'_j$ and $L(x, \succ_i) \subseteq L(x, \succ'_i)$ hold, $f(\succ) = x \in Q_i$ implies $f(\succ') = x$.*

Weak non-bossiness means that if alternative x is chosen for preference profile $\succ$ and x is in agent i's indifference set, for preference profile $\succ'$ in which only agent i's preference is changed such that x is not considered worse, x is also chosen. Note that Q_i is public information and agent i cannot report a preference that is inconsistent with it. Thus, since x is not considered worse, all of the elements in Q_i are not considered worse.

Standard non-bossiness, which was introduced by [15], means that by changing an agent's preference relation, she cannot change the outcome for the other agents without affecting her own outcome. Here, we use a weaker definition of non-bossiness; by changing agent i's preference in a restricted manner, i cannot change the outcome from one alternative in Q_i to another alternative in it. If an agent could do this, then she could change the outcome for the other agents without affecting her own outcome.

The following theorem holds:

Theorem 2. *Assume social choice function f is IR, SP and WNB. Then, the size of its possible multiagent choices, i.e., $|M(f)|$ is at most $n-1$ and its agent count, i.e., $C(M(f))$ is at most $2(n-1)$.*

To prove this theorem, we utilize a property called Maskin monotonicity [11].

Definition 7 (Maskin Monotonicity). *Social choice function f is Maskin Monotonic (MM) if for all $\succ, \succ' \in \mathcal{P}^n$, $f(\succ) = x$ and $L(x, \succ_i) \subseteq L(x, \succ'_i)$ for all $i \in N$ imply $f(\succ') = x$.*

In words, social choice function f is MM if alternative x, which is chosen for preference profile $\succ$ is also chosen for preference profile $\succ'$, where x is not considered worse for all agents. When the agents have strict preferences over X, strategy-proofness implies Maskin monotonicity [11]. On the other hand, in our model where agent i considers all alternatives in Q_i indifferent, strategy-proofness does not imply Maskin monotonicity. However, if we assume weak non-bossiness, strategy-proofness does imply Maskin monotonicity, and the following proposition holds.

Proposition 1. *If social choice function f is SP and WNB, then it satisfies Maskin monotonicity.*

Proof. To show that strategy-proofness and weak non-bossiness imply Maskin monotonicity, assume that $f(\succ) = x$. It is sufficient to show $f(\succ') = x$ when $\succ'$ satisfies $\succ_j = \succ'_j$ for any agent $j \in N_{-i}$ as well as $L(x, \succ_i) \subseteq L(x, \succ'_i)$ holds. Consider the following three cases:

case (1) $x \sim_i f(\succ')$: either $f(\succ') = x$ or $x \in Q_i$ holds. If $f(\succ') = x$, we are done. If $x \in Q_i$, we obtain $f(\succ') = x$ by weak non-bossiness.

case (2) $f(\succ') \succ_i x$: agent i has an incentive to declare $\succ_i'$ if her true preference is $\succ_i$. This fact violates the assumption that f is SP.

case (3) $x \succ_i f(\succ')$: we obtain $x \succ_i' f(\succ')$ from assumption $L(x, \succ_i) \subseteq L(x, \succ_i')$. Thus, agent i has an incentive to declare $\succ_i$ when her true preference is $\succ_i'$. This fact violates the assumption that f is SP.

As a result, $f(\succ) = f(\succ')$ holds. Thus, strategy-proofness and weak non-bossiness imply Maskin monotonicity.

For $x, x' \in X$, x dominates x' if $x \succeq_i x'$ holds for all $i \in N$ and there exists at least one agent $j \in N$ such that $x \succ_j x'$ holds.

The following proposition holds.

Proposition 2. *Assume that social choice function f is IR, SP, and WNB. When there exists alternative $x' \in X(f)$ that dominates alternative $x \in X(f)$ for $\succ$, $f(\succ) \neq x$ holds. In particular, if there exists alternative $x \in X(f)$ that dominates q, we have $f(\succ) \neq q$.*

Proof. For the sake of contradiction, we assume that there exists $\succ \in \mathcal{P}^n$ s.t. x' dominates x and $f(\succ) = x$. From $f(\succ) = x$ and individual rationality, for all $i \in S(x)$, $x \succ_i q$ holds. Also, from the assumption that x' dominates x, for all $i \in S(x') \cap S(x)$, $x' \succ_i x \succ_i q$ holds. Also, since we assume $x' \in X(f)$, there exists $\succ' \in \mathcal{P}^n$ s.t. $f(\succ') = x'$. From individual rationality, for all $i \in S(x')$, $x' \succ_i' q$ holds. Here, we consider $\succ''$ that satisfies the following conditions:

- If $i \in S(x) \cap S(x')$, $x' \succ_i'' x \succ_i'' q$.
- If $i \in S(x') \backslash S(x)$, $x' \succ_i'' q$ and $x \sim_i'' q$.
- For all $x'' \neq x, x'$, for all $i \in S(x'')$, $q \succ_i'' x''$.

$L(x, \succ_i) \subseteq L(x, \succ_i'')$ holds. Thus, from Maskin monotonicity, $f(\succ'') = x$ must hold. On the other hand, $L(x, \succ_i') \subseteq L(x, \succ_i'')$ holds for all $i \in N$. Thus, from Maskin monotonicity, $f(\succ'') = x'$ must hold. This is a contradiction.

Next, we introduce the classification of agents.

Definition 8 (Agent Types).

Decider: *Agent i is a decider if $|X_i(f)| \geq 2$ and $|M_i(f)| \geq 1$ hold, where at least two possible alternatives of mechanism f are her standard alternatives and at least one is a multiagent choice.*

Approver: *Agent i is an approver if $X_i(f) = M_i(f)$ and $|X_i(f)| = 1$ hold, i.e., exactly one possible alternative of mechanism f is her standard alternative and also a multiagent choice.*

Solitary: *Agent i is solitary if $|M_i(f)| = 0$ holds, where no possible multiagent choice of mechanism f is her standard alternative. There might exist alternative x such that $S(x) = \{i\}$ holds.*

In Example 1, agent 1 is a decider since $X_1(f) = \{x_3, x_4, x_5\}$, and x_3 and x_4 are multiagent choices. Agent 2 and 3 are approvers since $X_2(f) = M_2(f) = \{x_3\}$ and $X_3(f) = M_3(f) = \{x_4\}$.

The next proposition implies that if alternative $x \in X(f)$ is a multiagent choice, then there exists at most one decider who considers x her standard alternative. In Example 1, $X(f)$ includes two multiagent choices, i.e., x_3 and x_4. For both, only agent 1 is the decider of these alternatives.

Proposition 3. *Assume that agents $i, j \in N$ ($i \neq j$) are deciders. If there exists SP, IR, and WNB social choice function f, $X_i(f) \cap X_j(f) = \emptyset$ holds.*

Proof. For the sake of contradiction, we assume that agents 1 and 2 are deciders and $X_1(f) \cap X_2(f) \neq \emptyset$ holds. Consider the following two cases.
case (i) $|X_1(f) \cap X_2(f)| \geq 2$: Without loss of generality, assume $x, x' \in X_1(f) \cap X_2(f)$ $(x \neq x')$, $x \succ_1 x' \succ_1 q$ for agent 1, $x' \succ_2 x \succ_2 q$ for agent 2, and for any $x'' \in X \backslash \{x, x', q\}$, there exists agent $i \in S(x'')$ whose preference is $q \succ_i x''$. From individual rationality, $f(\succ) \neq x''$ holds. Furthermore, assume $x, x' \succeq_i q$ holds for any agent $i \in N \backslash \{1, 2\}$. Also, from Proposition 2, $f(\succ) \neq q$ holds. Without loss of generality, assume $f(\succ) = x$. Then, agent 2 has an incentive to declare $x' \succ'_2 q \succ'_2 x$ when her true preference is $\succ_2$. This is because when agent 2 declares $\succ'_2$, the social choice becomes x'. This contradicts our assumption that f is SP.
case (ii) $|X_1(f) \cap X_2(f)| = 1$: Without loss of generality, assume $x \in X_1(f) \cap X_2(f), x' \in X_1(f) \cap Q_2$, and $x'' \in Q_1 \cap X_2(f)$. We consider the following five cases of preferences for agents 1 and 2. For each case, we assume that for any $x''' \notin \{x, x', x'', q\}$, there exists agent $i \in S(x''')$ whose preference is $q \succ_i x'''$. We also assume $x, x', x'' \succeq_i q$ for agent $i \in N \backslash \{1, 2\}$. From individual rationality and Proposition 2, $f(\succ) \in \{x, x', x'', q\}$.

(1) agent 1: $x' \succ_1 q \sim_1 x'' \succ_1 x$,
agent 2: $x \succ_2 q \sim_2 x' \succ_2 x''$
(2) agent 1: $x' \succ_1 x \succ_1 q \sim_1 x''$,
agent 2: $x \succ_2 q \sim_2 x' \succ_2 x''$
(3) agent 1: $x \succ_1 x' \succ_1 q \sim_1 x''$,
agent 2: $x \succ_2 x'' \succ_2 q \sim_2 x'$
(4) agent 1: $x' \succ_1 x \succ_1 q \sim_1 x''$,
agent 2: $x \succ_2 x'' \succ_2 q \sim_2 x'$
(5) agent 1: $x' \succ_1 x \succ_1 q \sim_1 x''$,
agent 2: $x'' \succ_2 x \succ_2 q \sim_2 x'$

We show the social choice for each case.

(1) From individual rationality and Proposition 2, we have $f(\succ) = x'$.
(2) From (1) and Maskin monotonicity, $f(\succ) = x'$ holds $(*)$.
(3) From Proposition 2, we have $f(\succ) = x$ $(**)$.

(4) We have $f(\succ) = x'$ from the following three reasons: (a) If $f(\succ) = x$, it is contrary to strategy-proofness for agent 2 from $(*)$. (b) If $f(\succ) = x''$,it is contrary to $(**)$ from MM. (c) From Proposition 2, $f(\succ) \neq q$.

(5) From (4) and MM, $f(\succ) = x'$ holds. However, we can create four cases by replacing x' with x'', and agent 1 with agent 2 in cases from (1) to (4), which we call (1') to (4'). For example, in (1'), we set $x \succ_1 q \sim_1 x'' \succ_1 x'$ for agent 1 and $x'' \succ_2 q \sim_2 x' \succ_2 x$ for agent 2. Although $f(\succ) = x'$ has to continue holding in case (5), from (1') to (4'), we obtain $f(\succ) = x''$ in case (5). This is a contradiction.

Now, we are ready to prove Theorem 2.

Proof (Proof of Theorem 2). From Proposition 3, the number of deciders who consider each multiagent choice $x \in M(f)$ their standard alternative is at most one. Since there exists at least one approver for each $x \in M(f)$, we obtain $|M(f)| \leq n$. If $|M(f)| = n$, the number of approvers is n. However, this implies that the number of deciders is 0 and violates the assumption that there are n multiagent choices, since each alternative is a standard alternative only for a single approver. Thus, we derive $|M(f)| \leq n - 1$.

Next, we examine the agent count. If there exists no decider, then the agent count is at most n, where all approvers consider a single multiagent choice as their standard choice. If there exists at least one decider, then the agent count is at most the sum of the number of approvers and multiagent choices. This number is maximized when a single decider considers all $n - 1$ multiagent choices as her standard alternative, and each multiagent choice is considered a standard alternative by a single approver. In such a case, the agent count is $2(n - 1)$.

7 AC Mechanisms

We propose a class of IR, SP, and WNB social choice functions (mechanisms) that is optimal in terms of possible multiagent choices and agent count. We call such mechanisms *Approve and Choose (AC) mechanisms*. For presentation purposes, we first show the AC mechanism for 2 agents.

Mechanism 1 (AC Mechanism for 2 Agents). *We assume that there exist agents 1 and 2. Mechanism f chooses one decider and one approver in some way (which can be arbitrary) independently from the agents' declarations. Here, let assume agent 1 is a decider and agent 2 is an approver. $X(f)$ consists of $\{x^*\} \cup Y$. It determines alternative x^* as a multiagent choice which is a standard alternative for both agent 1 and 2, in some way independently from the agents' declarations. Thus, $x^* \in X_1(f) \cap X_2(f)$ holds. Also, it chooses set of alternatives $Y = \{q\} \cup \{x \in X_1(f) \colon |S(x)| = 1\}$, where Y contains q and all the alternatives are standard alternatives only to agent 1 in some way independently from the agents' declarations.*

1. *Agent* 2 *approves* x^* *if she prefers* x^* *over* q*. Otherwise, she vetoes* x^**.*
2. *Agent* 1 *selects the best alternative according to agent* 2*'s decision:*
   – *If agent* 2 *approves* x^**, then agent* 1 *selects the most preferred alternative in* $\{x^*\} \cup Y$*.*
   – *Otherwise, agent* 1 *selects the most preferred alternative in* Y*.*

We generalize this procedure from 2 agents to n agents. First, the mechanism chooses one decider in some way independently from the agents' declarations. Here, we assume agent 1 is a decider. The mechanism then determines a set X^* of multiagent choices in some way independently from the agents' declarations, such that for each $x^* \in X^*$, agent 1 is the decider and a subset of N_{-1} is approvers. Here, each approver has exactly one alternative in X^* and this choice is her standard alternative. In Example 1, the mechanism chooses X^* as $\{x_3, x_4\}$. For x_3, agent 1 is the decider and agent 2 is the approver. For x_4, agent 1 is the decider and agent 3 is the approver.

Also, the mechanism independently selects set of alternatives $Y = \{q\} \cup \{x \in X_1(f) \colon |S(x)| = 1\}$ from the agents' declarations. In Example 1, the mechanism chooses Y as $\{x_5, q\}$ and asks each approver in turn whether she approves of her multiagent choice. Let $\widehat{X}^* \subseteq X^*$ denote the set of alternatives that are approved by all own approvers. Then, the decider selects her most preferred alternative among $\widehat{X}^* \cup Y$.

In Example 1, assume $x_3 \succ_1 x_4 \succ_1 x_5 \succ_1 q$, $x_3 \succ_2 q$, and $x_4 \succ_3 q$. Then, x_3 and x_4 are approved by agents 2 and 3, respectively. Thus, $\widehat{X}^* = \{x_3, x_4\}$. Then, agent 1 chooses x_3 from $\widehat{X}^* \cup Y = \{x_3, x_4, x_5, q\}$. If the preference of agent 1 is $x_5 \succ_1 x_4 \succ_1 x_3 \succ_1 q$, then she chooses x_5.

The following theorem illustrates the property of the AC mechanism.

Theorem 3. *Any AC mechanism* f *is IR, SP, and WNB. Also,* $X(f) = X^* \cup Y$*. Furthermore, an instance of the AC mechanisms is optimal in terms of the number of multiagent choices and the agent count.*

Proof. Due to space limitations, we show the proof for the two agent case. Generalizing to the n agent case is rather straightforward (although verbose).

We first show that the AC mechanism is SP. For agent 2 (an approver), if she vetoes x^* even though $x^* \succ_2 q$, the social choice results in her indifference set. If she approves x^*, x^* might be selected. Thus, agent 2 has no incentive to misreport her preference. Agent 1 (a decider) does not have any incentive to misreport her preference, since she can select the best alternative from a set that is determined independently from her declaration. Next we show that the AC mechanism is IR. Obviously, it selects an alternative that is not worse than q as long as the agents truthfully declare their preferences. Then we show that the AC mechanism guarantees weak non-bossiness. Assume agent 1 chooses $y \in Y$ (which is in Q_2) when agent 2 approves x^*. Then if she vetoes x^* by declaring that she prefers y over x^*, agent 1 still chooses y. For agent 1, q is the only alternative in Q_1 that can be selected. If agent 1 chooses q, then under any preference of agent 1 in which q is not considered worse, q is also selected. x^*

can be chosen when both agents deem x^* best. Also, $x \in Y$ can be chosen if agent 1 deems x the best.

Finally, by selecting each multiagent choice so that it is considered a standard alternative by exactly one approver, the number of multiagent choices becomes $n-1$. The agent count becomes $2(n-1)$. From Theorem 2, this instance of AC mechanisms is optimal in terms of the number of multiagent choices and the agent count.

Since the agent count is bounded and there exists at most one decider for each alternative, there exists a tradeoff between the numbers of multiagent choices $|M(f)|$ and agents who consider $x \in M(f)$ their standard alternative $|S(x)|$; if $|S(x)|$ becomes large, $|M(f)|$ should become small. As shown in the proof of Theorem 3, $|M(f)|$ is maximized when $|S(x)| = 2$ holds for all $x \in M(f)$.

Unfortunately, the AC mechanism does not characterize all SP, IR, and WNB mechanisms, since we can consider the following (seemingly less attractive) mechanisms.

A Mechanism with Multiple Deciders: From Proposition 3, multiple deciders should not consider the same multiagent choice as their standard alternative. Assume agents are partitioned into several groups that are serialized. Each group contains exactly one decider and the approvers who consider the multiagent choices of the decider as standard alternatives. For the first group, we apply the above procedure of Mechanism 1. When the first group chooses q, the mechanism proceeds to the second group, and so on.

A Mechanism with No Decider: Assume all agents are solitary and serialized. The first agent is asked to choose her most preferred alternative. If her choice is q, the mechanism proceeds to the next agent, and so on.

8 AC Mechanism for Quasi-Linear Utilities

We can extend the idea of AC mechanisms and apply them to the cases where each agent's utility is quasi-linear.

An SC-EI problem in quasi-linear utilities is defined by a tuple (N, X, q, v, Q). The meanings of N, X, q, and Q are basically identical to SC-EI. Here, $v = (v_i)_{i \in N} \in V^n$ is a profile of the valuation functions, where V is the set of all possible valuation functions for an agent. Each valuation function $v_i : X \to \mathbb{R}$ returns the valuation of each alternative. We assume $v_i(x) = 0$ if $x \in Q_i$ holds. Mechanism $f = (g, p)$ consists of allocation function $g : V^n \to X$ and payment function $p : V^n \to \mathbb{R}^n$. When the declared profile of the valuation functions is $\widehat{v}$, the utility of agent i is defined as: $v_i(g(\widehat{v})) + p_i(\widehat{v})$, i.e., a quasi-linear utility function.

We require that the mechanism does not lose or earn money.

Definition 9 (Budget-Balanced). *A mechanism is* strongly budget-balanced *(SBB) if for any profile of valuations $v \in V^n$, $\sum_{i \in N} p_i(v) = 0$ holds.*

We say a mechanism is *weakly budget-balanced* (WBB) if for any profile of valuations $v \in V^n$, $\sum_{i \in N} p_i(v) \leq 0$ holds. The family of Groves mechanisms is a well-known representative class of efficient and SP social choice functions [8]. In our problem setting, agents can have a negative valuation for an allocation/choice. In such a setting, an instance of Groves mechanisms can satisfy IR, but not WBB. Another instance of Groves mechanisms can satisfy WBB, but not IR. Actually, no instance of a Groves mechanism is simultaneously IR and WBB [9]. We show that the modified AC mechanism guarantees SP, IR, and SBB.

Note that we do not require that WNB be a necessary condition in this section. Since an agent can affect the outcome through payments, it is common that a mechanism does not satisfy non-bossiness. For example, the Groves mechanisms do not satisfy non-bossiness.

8.1 Class of AC Mechanisms

In this section, we propose a class of AC mechanisms for quasi-linear utilities.

Mechanism 2 (AC Mechanism for Quasi-Linear Utilities)

1. *The mechanism selects one agent as a decider. WLOG, we assume agent* 1 *is the decider.*
2. *It chooses* X^* *and* Y *similar to an AC mechanism. In more detail, it selects* X^* *of multiagent choices such that for* $x^* \in X^*$, *agent* 1 *is a decider and subset* N_{-1} *is approvers. Each approver has exactly one standard alternative in* X^*. *It also independently chooses a set of alternatives* $Y = \{q\} \cup \{x \in X_1(f) \colon |S(x)| = 1\}$.
3. *Each approver* j *declares her (not necessarily true) valuation (denoted as* $\widehat{v}_j$*).*
4. *For each approver* j *of alternative* x, *her threshold price* $p_j(\widehat{v})$ *is determined. It can depend on the declaration of other approvers, but it must be independent from her own declaration.*
5. $x \in X^*$ *remains valid if* $\widehat{v}_j(x) \geq p_j(\widehat{v})$ *holds for each approver* j *of* x. *Let* $\widehat{X}^* \subseteq X^*$ *denote a set of valid alternatives. We also require that for each valid alternative* $x \in \widehat{X}^*$ *and for each approver* j *of* x, $p_j(\widehat{v})$ *is determined independently from the declarations of the approvers of other valid alternatives. In other words, if* $p_j(\widehat{v})$ *depends on* $\widehat{v}_k$, *then the alternative related to agent* k *is not approved.*
6. *For each* $x \in \widehat{X}^*$, $r(x) = \sum_{j \in S(x), j \neq 1} p_j(\widehat{v})$.
7. *Agent* 1 *(decider) chooses her most preferred alternative within* $\widehat{X}^* \cup Y$. *When* $x \in \widehat{X}^*$ *is chosen, agent* 1 *receives* $r(x)$, *and each approver* j *of* x *pays* $p_j(\widehat{v})$. *If agent* 1 *chooses* $x \in Y$, *no agent pays/receives anything.*

The following theorem illustrates the properties of the AC mechanism for quasi-linear utilities.

Theorem 4. *The AC mechanism for quasi-linear utilities is SP, IR, and SBB.*

Proof. We first show that the AC mechanism is SP. Agent 1 can choose her most preferred alternative in a fixed set. The set and the corresponding payments are determined independently from her own declaration. Thus, agent 1 has no incentive to misreport her valuation. For each approver j, her threshold price of the alternative related to agent j is determined independently from her own declaration. Also, her own declaration affects the payment of other valid alternatives only when her alternative is not approved. Thus, agent j has not incentive to misreport her valuation.

The AC mechanism is IR, since each approver accepts an alternative only if she can pay a threshold price. Also, the decider always has an option to choose q. Obviously, it satisfies SBB, since the sum of approvers' payments is transferred to the decider.

8.2 Instances of AC Mechanisms

We introduce the instances of AC mechanisms for quasi-linear utilities. The first is called a *fixed price mechanism*, in which a common fixed threshold price is used. To achieve good efficiency, the fixed threshold price must be determined appropriately. To do this, we need precise prior knowledge about the distribution of agent valuations. On the other hand, the second, which is called a *minimum-value k + 1-st price mechanism*, is more flexible; the threshold price is determined based on the valuations of other agents.

Mechanism 3 (Fixed Price Mechanism (FPM)). *For all $x \in X^*$ and for all approvers, we set constant c as a threshold price. Each alternative $x \in X^*$ remains valid if $\widehat{v}_i(x) \geq c$ holds for each approver. If $x \in X^*$ is chosen, the decider receives $(|S(x)| - 1) \times c$. Clearly, each threshold price of an approver is determined independently from all the approver declarations, including her own declaration.*

Here, we show an example how our FPM works.

Example 2. Consider the same situation as Example 1, but each agent has a valuation for an alternative. We assume that each agent has a profile of valuation functions for $(x_1, x_2, x_3, x_4, x_5, q)$:

- agent 1 : $v_1 = (-200, -100, 20, 250, 100, 0)$
- agent 2 : $v_2 = (250, 200, 150, 0, 0, 0)$
- agent 3 : $v_3 = (100, 100, 0, 50, 0, 0)$

Let's assume that agent 1 is selected as a decider. In this example, the mechanism selects $X^* = \{x_3, x_4\}$ and $Y = \{x_5, q\}$. It also determines a threshold price 100 for approvers. When each approver wants to travel with agent 1 by paying this threshold price, say, her valuation is not less than 100, she approves this offer. For x_3, agent 2 is an approver and for x_4, agent 3 is an approver. In this case, only agent 2 approves the offer, i.e., $\widehat{X}^* = \{x_3\}$. Thus, agent 1 selects her most preferred alternative from $Y \cup \widehat{X}^* = \{x_3, x_5, q\}$. For agent 1, when she selects x_3, her utility becomes $20 - (-100) = 120$ by receiving 100. On the other hand, if she selects x_5, her utility is 100. As a result, agent 1 selects x_3 and agents 1 and 2 are going to travel in Thailand.

Mechanism 4 (Minimum-Value $k+1$-st Price Mechanism (MPM)). *For each alternative $x \in X^*$, let $t(x)$ be the minimum of $\widehat{v}_j(x)$ for all approvers of x. Sort X^* in decreasing order of $t(x)$. Top $k \leq |X^*| - 1$ alternatives remain valid and constitute $\widehat{X}^*$. Let x^{k+1} denote the $k+1$-st alternative. For each valid alternative x, $p_j(\widehat{v})$ is determined to be $t(x^{k+1})$. We can assume the threshold price of an invalid alternative is equal to $t(x^k)$. It is clear that the threshold price of each approver is determined independently from her own declaration. Also, the threshold price of each approved alternative is determined independently from the declarations of the other approvers of valid alternatives.*

In the last of this section, we show an example of MPM.

Example 3. Consider the situation described in Example 2.

The mechanism selects $X^* = \{x_3, x_4\}$ as the same as above. We assume that each agent truthfully declares her valuation functions. Thus, we have decreasing order of $t(x)$: $150 > 50$. As a result, agent 2 is going to pay 50 for agent 1. If agent 1 selects x_3, her utility becomes $20 + 50 = 70$. However, for agent 1, her best alternative is x_5, since her utility is 100 by selecting x_5. As a result, agent 1 is going to travel in Thailand alone.

9 Conclusion

We investigated the IR and SP social choice functions for settings where the agents need to choose from a set of alternatives including status quo q. We first showed a negative result, i.e., IR and SP mechanism can choose at most one alternative besides q. To overcome this negative result, we introduced the SC-EI setting, where the indifference set of each agent is publicly known. We developed a class of IR and SP mechanisms that work in this setting called Approve and Choose (AC) mechanisms, which can be optimal in terms of possible multiagent choices/agent counts.

Our future work will extend the AC mechanisms in a setting with monetary transfers and experimentally evaluate our mechanisms in various application domains.

Acknowledgements. This work was partially supported by JSPS KAKENHI Grant Numbers 24220003 and 15H02751, and JSPS Program for Advancing Strategic International Networks to Accelerate the Circulation of Talented Researchers.

References

1. Bachrach, Y., Elkind, E., Faliszewski, P.: Coalitional voting manipulation: a game-theoretic perspective. In: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI 2011), pp. 49–54 (2011)
2. Barberá, S., Ehlers, L.: Free triples, large indifference classes and the majority rule. Soc. Choice Welf. **37**(4), 559–574 (2011)
3. Blau, J.H., Deb, R.: Social decision functions and the veto. Econometrica **45**(4), 871–897 (1977)
4. Brandt, F., Geist, C.: Finding strategyproof social choice functions via SAT solving. In: Proceedings of the 14th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2014), pp. 1193–1200 (2014)
5. Darmann, A., Elkind, E., Kurz, S., Lang, J., Schauer, J., Woeginger, G.: Group activity selection problem. In: Goldberg, P.W. (ed.) WINE 2012. LNCS, vol. 7695, pp. 156–169. Springer, Heidelberg (2012)
6. Faltings, B.V.: A budget-balanced, incentive-compatible scheme for social choice. In: Faratin, P., Rodríguez-Aguilar, J.-A. (eds.) AMEC 2004. LNCS (LNAI), vol. 3435, pp. 30–43. Springer, Heidelberg (2006)
7. Gibbard, A.: Manipulation of voting schemes: a general result. Econometrica **41**(4), 587–601 (1973)
8. Groves, T.: Incentives in teams. Econometrica **41**(4), 617–631 (1973)
9. Guo, M., Shen, H., Todo, T., Sakurai, Y., Yokoo, M.: Social decision with minimal efficiency loss: an automated mechanism design approach. In: Proceedings of the 15th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2015), pp. 347–355 (2015)
10. Lu, T., Boutilier, C.: Multi-winner social choice with incomplete preferences. In: Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013), pp. 263–270 (2013)
11. Maskin, E.: Nash equilibrium and welfare optimality. Rev. Econ. Stud. **66**(1), 23–38 (1999)
12. Peleg, B.: Consistent voting systems. Econometrica **46**(1), 153–161 (1978)
13. Sato, S.: Strategy-proof social choice with exogenous indifference classes. Math. Soc. Sci. **57**(1), 48–57 (2009)
14. Satterthwaite, M.A.: Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. J. Econ. Theory **10**, 187–217 (1975)
15. Satterthwaite, M.A., Sonnenschein, H.F.: Strategy-proof allocation mechanisms at differentiable points. Rev. Econ. Stud. **48**(4), 587–597 (1981)