Compromise and strategyproof social choice under uncertainty

A special case

Author : Matteo Pianella

Thesis Director : François Maniquet, professor of economics at UCLouvain

Thesis Reader : Edoardo Grillo, assistant professor of economics at Collegio Carlo Alberto

Academic Year 2019 - 2020

In order to obtain the joint degree:

Master 120 en Sciences économiques, Orientation générale, Finalité approfondie (UCL/UNamur)

and

Dottore magistrale in Economics and Political Science (UNIMI)

Université catholique de Louvain - Economics School of Louvain - Place Montesquieu 3, 1348 Louvain-La-Neuve
Université de Namur – Economics School of Namur – Rempart de la Vierge 8, 5000 Namur
Università degli Studi di Milano • Via Festa del perdono 7 • 20122 Milano / Italy

Acknowledgements

This work has benefited form the support of many people. Prof. Maniquet has been of particular help in many stages of this work. His technical suggestions have been essential to guide this work towards a better outcome. His critical remarks pushed me to growth intellectually. I want to thank prof. Grillo to have accepted to read and comment on my work with flexibility and availability. A special thank goes to Ms. Carette, Ms. Gilot, Ms. Cottone and Mr. Parenti, without whom I would now be lost in the midst of some bureaucratic intrigue. Last but not least, a big thanks to my family and friends whose presence gave me energy and motivation along the way.

Contents

Chapter 1

Introduction

An important problem of any group decision is to aggregate the preferences of the individuals in such a way that they cannot manipulate this decision to their advantage by misreporting their preferences. Inside social choice theory, any rule that has this property of non manipulability is called a *strategyproof* social choice function (also SCF). The problem of finding domains of preferences that admit non trivial strategyproof SCF is an old topic of social choice theory.[1]

The results obtained independently by Allan Gibbard and Mark Satterthwaite (Gibbard 1973, Satterthwaite 1975) sign the beginning of this enquiry. They established that in the full domain of preferences with more than two alternatives the social choice function must be dictatorial, in what is known as the Gibbard-Satterthwaite theorem. This impossibility results has been followed by many other impossibility results in domains subject to different kinds of constraints. Overall, this literature has established that the requirement of strategyproofness is difficult to meet and it usually involves trade-offs with other important features of the SCFs, like efficiency. However, there are some domains, most famously the domain of single peaked preferences (Moulin 1980), that admit strategyproof functions with very interesting properties.

This 45 years of research in this field seemed to have covered most of the possible interesting domains of preferences; for sure the most important domains seemed to have been addressed. Instead, a new paper by Bahel and Sprumont 2020 (also BS) shows that this is not the case when they reveal that the domain of *subjective expected utility* (also SEU) preferences has been overlooked. The domain of SEU preferences constitutes an important gap in the literature and this lack is even more surprising given that the adjacent domain of group decisions under risk has received attention since when the research program on strategyproof SCFs started (Zeckhauser 1973).

Bahel and Sprumont 2020 analyses a case in which a group decision must be taken in conditions of uncertainty where agents can express their preferences among different social acts. A social act is a prospect that associates an outcome to every possible realization of the world and agents' preferences over these social acts are represented by their subjective expected utility. The authors are able to characterize the class of SCFs that are strategyproof and that respect a weak unanimity axiom.

Even though this class of functions has interesting features, it has drawbacks, too. One of the main problems is that they only chose social acts that in any state of the world select someone's top outcome. In fact, there are situations in which the agents would prefer a

[1]For a very good although somewhat old survey of this literature see Barberà 2011

compromise act to the ones that give to them dictatorial power in some state of the world; the functions in BS do never select the compromise act. This problem has been explicitly recognized by the authors and this is the problem that we want to address in this study.
Therefore, our main question is the identification of strategyproof SCFs that compromise. Besides the normative appeal of a compromise result of this kind, we conducted this analysis with the idea that by relaxing the requirement of unanimity and introducing another axiom compatible with strategyproofness we would be able to find a new class of functions in this domain that could display interesting properties.
We analyse this question in the special case of two agents, three outcomes and two states of the world. By imposing a particular axiom that requires the SCF to compromise in particular profiles of preferences we are able to produce a full characterization of all the strategyproof SCFs that compromise in some profiles of preferences. We will be able to show that these function exhibit interesting properties not present in the functions characterized in BS and we will make clear that these functions constitute a different way of producing social choices under uncertainty.
The study is divided in two main chapters. Chapter 2 is almost entirely devoted to a detailed presentation of Bahel and Sprumont 2020. At the end of the Chapter we will refer to the most relevant literature, but we maintain this last section brief to avoid repetition with BS. In Chapter 3 we present our results. After specifying our model, we derive some new implications of strategyproofness when there is a unilateral deviations in subjective beliefs. These new findings are necessary to derive our main result, that is contained in Theorem 1. For reason of readability the proofs are relegated to the Appendix.
We close Chapter 3 with a comparison between the function that we have characterized and the functions in BS on three different dimensions. First, we argue that these classes of functions define two different ways of selecting social acts and we do so by pointing to the different properties that they respect. Then, we compare the efficiency property of one of our functions with one in BS. We close the section by discussing whether these two classes of functions admit group-strategyproof SCFs. We conclude the study by discussing some ideas for a possible extension of these results to a more general model.

Chapter 2

Review of the literature

2.1 Bahel and Sprumont (2020)

This entire work is motivated by Bahel and Sprumont 2020. The model that we will study in section 3 is indeed a particular case of the model in this paper. Therefore, some of the results derived here (the three lemmas on valuations presented in the next section) will also apply to section 3 and they will be extensively used there.
The main result of this paper is the characterization of the strategyproof and unanimous social choice functions in the domain of SEU preferences. We will present the model and the results in details in the next two sections.

2.1.1 The model

The paper models a situation in which a group decision among different uncertain prospects must be taken. There is a set set $N = 1, 2, ..., n$ of at least two agents. There are different states of the world that can realize. We call $\Omega = \{\omega_1, ..., \omega_m\}$, with $m \geq 2$ the set of states of the world (also called states of nature). Each state of the world is a "description of the world so complete that, if true and known, the consequence of every actions would be true and known"(Arrow 1971, p. 45). An event is any set of states of nature.
There is a set $X = \{a, b, ...\}$, of cardinality at least 3, of possible outcomes. These outcome are descriptions of what can happen to the agents and they are *public goods* in the sense that they are consequences that affects every agent. There are no private goods like money in this model.
Each agent has a subjective belief about the probability that a particular event will realize. In particular, consider the set of possible events 2^Ω, which is a σ-algebra of subsets of Ω. Agent i's beliefs are represented by a probability measure p_i over 2^Ω. Namely, p_i is a map $p_i : 2^\Omega \to [0, 1]$ such that : (a) $p_i(\emptyset) = 0$ and $p_i(\Omega) = 1$; (b) $p_i(A) \geq 0$, for all $A \in 2^\Omega$; (c) $p_i(\bigcup_q A_q) = \sum_q p_i(A_q)$ for any countable collections of disjoint sets $A_q \in 2^\Omega$. The fact that the beliefs are subjective means that each agent can assign a different probability to the same event.
Agents can also have different preferences over the outcomes. These preferences are represented by a valuation function v_i, that is a map $v_i : X \to \mathbb{R}$, normalized between 0 and 1. This map is state-independent, i.e. the valuation does not depend on which state the outcome is selected.[1]

[1]If the v_i depended on the state that is actual, this would contradict postulate 3 in Savage's axiomatization (Savage 1954). The results in Bahel and Sprumont 2020, as well as our results, do not apply to cases in

We call *social act* a map that associates to each state of the world an outcome. For instance, we write (a, b, c) to represent the social act that associates to state ω_1 the outcome a, to state ω_2 the outcome b and to state ω_3 the outcome c. In the model that is relevant to this study, the set of possible acts is the Cartesian product X^Ω.
Bahel and Sprumont 2020 also presents a model, that they call Constrained model, in which the set of possible outcomes X is contingent to the state of nature in which the outcome is selected. For instance, it is possible that under ω_1 we can select either a, b and c, but under ω_2 c is not anymore available, and only a or b can be chosen. We will not present this model here, since the results that we will develop in chapter 3 and 4 concern a situation in which the set of social acts is a Cartesian product. It is enough to notice that the case of a Cartesian product is a particular case of the Constrained model. Bahel and Sprumont 2020 derives the proofs of the Constrained model by minor adjustments on the arguments for the case with the Cartesian product.
One of the main assumptions of Bahel and Sprumont 2020 is that the agents' preferences over social acts are of the *subjective expected utility* type. This means that for any agent i, for any preference over acts $\succeq_i$ for any two acts $f, g \in X^\Omega$ there exists a pair (v_i, p_i) such that

$$f \succeq_i g \iff E^{p_i}_{v_i}(f) \geq E^{p_i}_{v_i}(g)$$

with, for any act h, $E^{p_i}_{v_i}(h) := \sum_{\omega \in \Omega} p_i(\omega) v_i(h(\omega))$ and $h(\omega)$ the outcome selected by act h at state ω. We say that the pair (v_i, p_i) represents the preference $\succeq_i$.
It must be noticed that the same preference $\succeq_i$ can be represented by more that two pairs of valuations and beliefs.[2] Nevertheless, if (v_i, p_i) and (w_i, q_i) both represent $\succeq_i$ then: (a) for each $x, y \in X$, $v_i(x) \geq v_i(y) \iff w_i(x) \geq w_i(y)$ and (b) for each $A, B \in 2^\Omega$, $p_i(A) \geq p_i(B) \iff q_i(A) \geq q_i(B)$. Namely, if (v_i, p_i) and (w_i, q_i) both represent $\succeq_i$, they generate the same order over valuations and over beliefs.
This model is limited by the assumption that $\succeq_i$ is a strict linear ordering over acts. This implies that for any pair (v_i, p_i) both the valuation function v_i and the probability p_i are injective functions (the case of p_i is shown at note 2 of the Appendix to Bahel and Sprumont 2020, and the case of v_i is straightforward).
Call $\mathcal{V}$ the set of normalized injective valuation functions v_i and $\mathcal{P}$ the set of possible injective probability measures over the set of possible events 2^Ω. Then the domain $\mathcal{D}$ of strict SEU preferences is

$$\mathcal{D} = \{(v_i, p_i) \in \mathcal{V} \times \mathcal{P} : E^{p_i}_{v_i}(f) \neq E^{p_i}_{v_i}(g) \text{ for all } f, g \in X^\Omega \text{ such that } f \neq g\}.$$

We denote with (v, p) the ordered list $((v_1, p_1), ..., (v_N, p_N)) \in \mathcal{D}^N$.

The group choice consists in selecting a social act in a way that depends on the profile of individual preferences. Then, it is natural to model this selection though a function, called *social choice function* (SCF) that is a map $\varphi : \mathcal{D}^N \to X^\Omega$. This is the mathematical object at the center of the analysis.
An SCF is *strategy-proof* if for all $(v, p) \in \mathcal{D}^N$ and for all i and $(v'_i, p'_i) \in \mathcal{D}$, then

$$E^{p_i}_{v_i}(\varphi(v, p)) \geq E^{p_i}_{v_i}(\varphi((v'_i, v_{-i}), (p'_i, p_{-i})))$$

which the valuation is state-dependent. For a philosophical discussion of the problem risen by postulate 3, see the correspondence between Robert Aumann and Leonard Savage that can be found in Drèze 1990. For a survey on state-depend preferences look at Drèze and Rustichini 2004.

[2] See Haller 1985 for a discussion of this point.

with $(v'_i, v_{-i}) = (v_1, ..., v_{i-1}, v'_i, v_{i+1}, ..., v_N)$ and $(p'_i, p_{-i}) = (p_1, ..., p_{i-1}, p'_i, p_{i+1}, ..., p_N)$.
When an SCF is strategyproof, no agent can strictly gain by misreporting his valuation function or his beliefs.
A first consequence that strategyproofness imposes on the SCFs is helpful to simplify the notation. In particular, we know that each profile of preferences $(\succeq_1, ..., \succeq_n)$ can be represented by more than one list of valuations and beliefs. Consider two lists $(v, p) \in \mathcal{D}^N$ and $(w, q) \in \mathcal{D}^N$ that represent the same profile of preferences $(\succeq_1, ..., \succeq_n)$. If the SCF φ is strategyproof, then $\varphi(v, p) = \varphi(w, q)$. This allows us to refer from now on to any pair $(v, p) \in \mathcal{D}^N$ as a *preference profile*. Moreover, we call any $(v_1, ..., v_n) \in \mathcal{V}^N$ a *valuation profile* and any $(p_1, ..., p_n) \in \mathcal{P}^N$ a *belief profile*.
The other axiom that Bahel and Sprumont 2020 imposes on SCF is unanimity. An SCF is *unanimous* if for all $(v, p) \in \mathcal{D}^N$ and all $f \in X^\Omega$,

$$[E^{p_i}_{v_i}(f) \geq E^{p_i}_{v_i}(g), \text{ for all } i \in N \text{ and for all } g \in X^\Omega] \Rightarrow [\varphi(v, p) = f].$$

This requirement of unanimity is very weak. We could compare it with other two efficiency requirements used in cases of risk or uncertainty, namely ex-ante Pareto and ex-post Pareto. The *ex-post Pareto* condition requires the social choice to be efficient once the uncertainty is resolved. Namely, if there exist $a, b \in X$ such that $v_i(a) > v_i(b)$ for all $i \in N$, then an ex-post Pareto social choice function will never choose b at any state of the world. Instead, the *ex-ante Pareto* condition imposes the efficiency requirement on the chosen acts, before the uncertainty is resolved. Formally, if there exist $f, g \in X^\Omega$ such that $E^{p_i}_{v_i}(f) > E^{p_i}_{v_i}(g)$ for all $i \in N$, then an ex-ante Pareto social choice function will never select g. The ex-ante condition is stronger than the ex-post condition, that is stronger than our unanimity requirement.[3]

Claim 2.1.1. *Ex-post Pareto is stronger than unanimity.*

Proof. Consider a profile (v, p) in which all agents agree that $a \in X$ is better than any other outcome in X, and $\varphi(v, p) = f$. Assume that $f \neq (a, ..., a)$, namely there is a state of the world $\omega \in \Omega$ in which φ selects an outcome $b \in X$, with $b \neq a$. Since a is strictly preferred to b by any agent and b is selected in one state, this contradicts ex-post Pareto.
To show that ex-post Pareto is strictly stronger than unanimity, take any situation in which the agents do not agree on their first ranked outcome, but they all agree that $b \in X$ is better than $c \in X$. Unanimity does not impose any restriction on the outcome selected in this case, while ex-post Pareto requires that c is never selected by the SCF. Therefore, the SCF may be unanimous without being ex-post Pareto efficient. □

Since our model requires that the set of possible acts is the Cartesian product X^Ω, the unanimity axiom simply imposes that when all agents agree on their best outcome, then the SCF must select this outcome at any state of the world. Because it is so weak, the authors call this axiom an auxiliary requirement, a technical restriction that is imposed with the primary aim of identifying a class of strategyproof SCFs. Whether this axiom, although very weak, is something more than a technical restriction will be discussed at the end of this section and again at the end of section 3.

[3] For the first implication see, for instance, Hylland 1980, p.16.

2.1.2 Results

The proof of the characterization result (see online Appendix to Bahel and Sprumont 2020) is long and requires several steps. As preliminary step to the proof, the authors derive three implications that the requirement of strategyproofness imposes on the SCFs in this domain. We call them here *lemmas on tastes* to underline the fact that these are results that primarily concern what a strategyproof SCF can do when the valuations over outcomes change, holding fixed the beliefs. It follows an exposition of the three lemmas. Throughout the exposition of the lemmas φ denotes a strategyproof SCF.

Let us write $\varphi(v, p; \omega)$ to denote the outcome that φ selects at state ω under the preference profile (v, p). For each $f \in X^\Omega$ and $x \in X$, let us denote with f^x the event under which outcome x is selected by the act f. In particular, $\varphi(v, p)^x := \{\omega \in \Omega : \varphi(v, p; \omega) = x\}$.

Monotonicity.
Let $i \in N$, $a \in X$, and let $(v, p), (w, p) \in \mathcal{V}^N$ be such that $v_i(a) > w_i(a), v_i(x) = w_i(x)$ for all $x \neq a$, and $v_{-i} = w_{-i}$. If $\varphi(v, p) \neq \varphi(w, p)$, then $p_i(\varphi^a(v, p) > p_i(\varphi^a(w, p))$.

In words, if the SCF selects a different act at two profiles such that one agent increases his valuation of one outcome in the first profile compared to the second (all else equal), then this agent's subjective belief that this outcome will be selected at the first profile must increase too.

Let us denote with $v_i \simeq w_i$ the fact that the valuations v_i and w_i generate the same ordering over outcomes, namely $(v_i(x) - v_i(y))(w_i(x) - w_i(y)) > 0$ for all $x, y \in X$. By abusing notation, we write $w \simeq v$ if $v_i \simeq w_i$ for all $i \in N$.

Ordinality.
If $(v, p), (w, p) \in \mathcal{D}^N$ and $v \simeq w$, then $\varphi(v, p) = \varphi(w, p)$.

In words, a strategyproof SCF must select the same act at two profiles with the same beliefs and the same order over outcomes.

Given $v_i \in \mathcal{V}$ we say that two outcomes $a, b \in X$ are *adjacent in v_i*, if no $x \neq a, b$ has utility $v_i(x)$ between $v_i(a)$ and $v_i(b)$. We say that w_i *obtains by permuting the utilities of a and b in v_i* if $w_i(a) = v_i(b)$, $w_i(b) = v_i(a)$, and $w_i(x) = v_i(x)$ for $x \neq a, b$.

Permutation invariance.
Let $(v, p) \in \mathcal{D}^N$, and let $a, b \in X$ be adjacent in v_i. If w_i obtains by permuting the utilities of a and b in v_i, then $\varphi^x(w_i, v_{-i}) = \varphi^x(v)$ for all $x \neq a, b$.

In words, a permutation of adjacent outcomes does not change the events at which the remaining outcomes are selected.

Since all the previous three results derive by simply assuming that SCF is strategyproof, we will use these lemmas in our proofs of the compromise result of chapter 3.
The first real step towards the characterization of unanimous and strategyproof rules comes with the proof of the *Top Selection Lemma.* This mainly states two things: (a) a strate-

gyproof and unanimous SCF must select in all state of the world someone's top alternative; (b) the event at which one's favourite outcome is selected varies only with beliefs; the valuations of non top outcomes do not matter for the final social choice. Point (b) is saying that the SCFs that the authors are going to characterize use a very limited amount of information in terms of agent's tastes. Even more important for our study is that, by point (a), any kind of compromise act is excluded, even in those situations in which this seems the natural solution.
To show this last point we use the example provided in BS. Take the case of three outcomes $\{a, b, c\}$, two states of the world and two agents A and B with valuations $v_A(a) = 1, v_A(b) = 0.99, v_A(c) = 0$ and $v_B(a) = 0, v_B(b) = 0.99, v_B(c) = 1$. The two agents strongly disagree on the valuations of alternatives a and b, while they agree that alternative b is very good. In this case the natural solution seems to be the act (b, b). Instead, point (b) of the Top Selection Lemma is limiting our choice to the four acts $(a, a), (a, c), (c, a), (c, c)$. Trying to reconcile this kind of situations with strategyproofness is a central motivation of our study.
The proof of the Top Selection Lemma is by induction on the unanimity axiom. In particular, they start with the situation in which an outcome ranked first by all agents. By unanimity, the SCF must select only the top outcome in this situation. Then they generalize this statement to all other possible valuation profiles using induction. Therefore, the introduction of unanimity reduces the class of strategyproof functions under the constraints imposed by the Top Selection Lemma. Then, we may get the intuition that relaxing this axiom, however weak it may be, could enlarge the class of SCFs in an interesting way. It is also interesting to notice that the unanimity axiom is only used in the Top Selection Lemma; it has no other implications throughout the proof.
Another important feature is that the beliefs are held fixed for the entire proof of the Top Selection Lemma. This is important for us because, due to this technical reason, the study of the implications of strategyproofness alone in this domain are not completely developed in Bahel and Sprumont 2020. The above three lemmas on tastes only concern a deviation in valuations and leave open the question of what can happen when a deviation in beliefs alone can take place or when a deviation in both beliefs and tastes can take place. This will be the subject of the analysis of section 3.2.
The Top Selection Lemma allows the analysis to focus on what they call the *selection functions*. This are functions that map events to agents and identify in which events a particular agent's top outcome will be selected. They can be seen as functions that allocates *rights over states of the world* to the agents, with which they have the right to decide which outcome to select if a particular state in the event they have been assigned realises. Once the Top Selection Lemma is established, the only thing that remains to establish is the particular shape of these selection rules. It is obvious that any SCF generated by a selection rule respects unanimity. It has still to be established which of those SCFs respect stategyproofness. The Characterization Theorem is saying that these functions coincide with the class that the authors call *locally bilateral top selection*.

Characterization Theorem *An SCF $\varphi : \mathcal{D}^N \to X^\Omega$ is strategyproof and unanimous if and only if it is a locally bilateral top selection .*

We conclude this section by describing these functions.
A first important element of these SCFs is that they are generated by the composition of "sub rules". In particular, taken a partition of Ω, for each event in that partition we can define a selection rule that assigns the states of nature in that event to the agents in one of

three possible ways that we are going to describe. The final social selection is generated by aggregating the selection operated in each event of the partition. Since the SEU preferences are separable (i.e the value of a social act is the sum of the values of its components), then if the marginal choice at each event is strategyproof, the final selection will be strategyproof, too.[4]

Take any partition of Ω. For each event inside this partition the control rights over states of nature is decided by one of the following "sub-rules".

Constant assignment. The assignment is fixed exogenously; agent's beliefs over states of nature does not influence the final assignment.

Bilaterally dictatorial assignment. Take any two agents and any proper covering of the event that we are considering. One of the two agents (the dictator) chooses one set of states in this covering and the other agent is assigned the complementary set. In this case the belief of the dictator directly influences the assignment, while the belief of the other agent is not relevant.

Bilaterally consensual assignment. Take two agents and consider a partition of the event that we are considering. Provisionally assign a set in that partition to one agent and the complementary set to the other agent. If both agents prefer the event assigned to the other agent then revert the assignment, otherwise maintain the initial assignment. In this case both agents' beliefs matter for the final assignment.

The description of these rules makes evident that assignment of the control right for each state of the world may only depend at most on the beliefs of two agents. This is why these SCFs are called *bilateral*, and it is an unexpected characteristic feature of the unanimous and strategyproof social choice functions in this domain.

An example may be helpful to better understand what kind of social choice functions can be generated by these rules.

Example. Take a situation with five states of nature $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$ and three agents, $N = \{A, B, C\}$.

Consider the following partition $P = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4, \omega_5\}\}$. An example of a locally bilateral top selection SCF is: event $\Omega' = \{\omega_1\}$ is assigned to agent A by a constant assignment; event $\Omega'' = \{\omega_2, \omega_3\}$ is assigned to agents B and C using a bilaterally dictatorial assignment, with B the dictator; event $\Omega''' = \{\omega_4, \omega_5\}$ is assigned to agents A and C by a bilaterally consensual assignment.

Under this rules, no matter the agents' beliefs, A's top outcome will be selected at state ω_1. The selection at event Ω'' depends on B's beliefs. Since the only proper covering of this event is $Q = \{\{\omega_1\}, \{\omega_2\}\}$, then if B thinks ω_2 is more likely than ω_3, B's top will be selected at ω_2, and C's top at ω_3. If B believes that ω_3 is more likely than ω_2, then the assignment is reversed.

At event Ω''' the selection depends on both A's and C's beliefs. There are two partitions of this event, namely $R = \{\{\omega_4\}, \{\omega_5\}\}$ and $R' = \{\{\omega_1, \omega_2\}\}$. If we take R' as our partition, then this rule would coincide with the constant assignment that gives both states to one

[4] Sprumont 1995, for instance.

agent. Then, take R and assume that we provisionally assign A to ω_4 and C to ω_5. In this case, if C believes that ω_4 is more likely than ω_5 *and* A believes that ω_5 is more likely than ω_4, then A's top will be chosen at ω_5 and C's top at ω_4. Otherwise, we reverse the selection.

Even though, due to the Top Selection Lemma, the amount of information concerning individual tastes that these SCFs can take into account is very limited, these functions is able to extract a fair amount of information on agents' beliefs. In particular, the bilaterally dictatorial rule may give us some detailed information about the beliefs of one agent and - when the state space allow for a partition of cardinality at least as big as the number of agents - we may extract some information for all agents. With the bilaterally consensual rules, the kind of information is different in that we can extrapolate differences in beliefs between agents.

It is important to realize that these social choice functions are very similar to the random dictatorship in Hylland 1980. The main difference is that, while the random dictatorships assign the control rights over states based on the realization of some event that have an objective probability, in our case these events are uncertain and are not associated with an objective probability anymore. The locally bilateral top selection rules are more complex in the way they attribute dictatorship, and some of them (bilaterally consensual) can welfare improve over random dictatorship. Nevertheless, they rise the same concerns risen by random dictatorship.

No matter how we build the SCF, in any realization of the world there will always be a dictator that will impose his decision over the others. This is obviously a normative problem ex-post, but it is also a problem ex-ante. Take any situation in which a large group of people with a large array of preferences must select an uncertain prospect that may contain many possible outcomes. If the number of possible state upon which these outcomes can be made contingent is big not big, then most of people's preferences will remain unheard. In addition to this, as we have already pointed out (see the above discussion of the Top Selection Lemma), these SFCs do not allow for any compromise act, even when this may seem the natural choice.

We end this section by suggesting another similarity between these strategyproof functions and those that emerge in the context of risk. As the random dictatorial rules, all the locally bilateral top selection rules are ex post Pareto. This comes immediately from the Top Selection Lemma that requires the SCF to select someone's top outcome in all states of nature. At the same time, it is a direct consequence of the Characterization Theorem that when the efficiency conditions is strengthened to ex-ante Pareto, then the only strategyproof functions respecting that criteria is the dictatorial one: an agent's top alternative is selected in all states of nature. This comes form the fact that for any other SCFs characterized by BS, it is always possible to find a profile of preferences in which both agents prefer a compromise result to the act that selects one of their top alternatives in all states of nature. This is exactly what happen in the context of risk: when ex ante Pareto is required, the the only strategyproof SCF is the dictatorial one.

2.2 Other related literature

The most related antecedent to BS is that strand of literature that has focused on the study of situations of voting under risk. A pioneering contribution in this direction is Zeckhauser 1973. In this paper, agents have to vote on lotteries of different possible outcomes. They have preferences over these outcomes and agent's preferences over lotteries are assumed to respect the Von Neumann-Morgenstern axioms. As a consequence, agent's preferences over lotteries are represented by their expected value. Another seminal contribution is Gibbard 1977. The approach here is different in that the outcomes are lotteries and the preferences are rankings over these outcomes. Besides the elegant characterization of the strategyproof rules, the author proves that there is a large trade-off between strategyproofness and efficiency in this domain. As also shown later by Hylland 1980, ex-ante efficiency shrinks the class of strategyproof functions to the dictatorial rule, while ex-post efficiency limits these functions to random dictatorial rules.

BS presents an extensive list of related work that we avoid to repeat here. Before closing this section it is worth mentioning that the previous efficiency requirements have also received attention from a normative standpoint. In particular, Mongin 2016 has criticized these two efficiency requirements. He argues that both these axioms imply what he calls spurious unanimity: cases in which the agents agree to prefer one social act over another, but they disagree on the underlying reasons for these preferences. If we take as underlying reasons the beliefs over states of the world and the valuations over outcome, then this evidently poses a problem for ex-ante efficiency when agents have SEU preferences. Indeed, the same act can be preferred to another by two agents that have opposite beliefs and valuations.

The ex-post Principle does not immediately fall in this problem, but the paper raises other important concern. Ex-post Pareto avoids the problem risen by the ex-ante principle by applying the Pareto principle at the level of the outcomes. Mongin 2016 argues that this does not solve the problem because the outcomes are really fixed and complete descriptions of a set of consequences, but "[t]hey should rather be seen as aggregates of further uncertain consequences, each of which could be analysed in turn, and so on ad infinitum"(Mongin 2016, p.529). This creates an inconsistency for the ex-post principle that loses his efficacy and it becomes of no possible use.

We have seen in the previous section that the unanimity requirement in BS is weaker than ex-post Pareto, therefore the first critique does not apply in this case. Nevertheless, a version of the second critique is still valid. Unanimity requires that when people agree on their best outcome than the SCF must select that outcome at all states. If we accept the description of outcomes as aggregates of further uncertain consequences, then even the weak unanimity requirement in BS would be of poor normative meaning.

Chapter 3

The compromise result

3.1 The model

Our model is a particular case of the model in Bahel and Sprumont 2020 presented in section 2.2.1. We take the smallest cardinality for the set of agents, outcomes and states of the world such that the results in Bahel and Sprumont 2020 still hold. The three lemmas on tastes presented in section 2.1.2 still apply to any strategyproof SCF in our environment and will be extensively used throughout the analysis.
In particular, there are two agents, that we call agent 1 and agent 2. Call $X = \{a, b, c\}$ the set of possible outcomes and $\Omega = \{\omega_1, \omega_2\}$ the set of states of the world. In this limited case there are only nine possible social acts:

$$X^\Omega = \{(a, a), (a, b), (a, c), (b, a), (b, b), (b, c), (c, a), (c, b), (c, c)\}.$$

To each agent is associated a valuation function v_i normalized between zero and one, and a subjective belief p_i, with $i = 1, 2$. Each agent i has strict preferences of the SEU type over the above 9 acts. Namely, given two acts $f, g \in X^\Omega$

$$f \succ_i g \iff p_i(\omega_1)v_i(f(\omega_1)) + p_i(\omega_2)v_i(f(\omega_2)) > p_i(\omega_1)v_i(g(\omega_1)) + p_i(\omega_2)v_i(g(\omega_2)),$$

with $f(\omega)$ and $g(\omega)$ the outcome selected respectively by act f and g at state ω. Recall that when the preferences over acts are strict, then the valuation function and the subjective belief are injective functions (see note 2 to Appendix to Bahel and Sprumont 2020).
Call $\mathcal{L}$ the domain of SEU preferences restricted to the case of two states of nature and three outcomes:

$$\mathcal{L} = \{(v_i, p_i) \in \mathcal{V} \times \mathcal{P} : E^{p_i}_{v_i}(f) \neq E^{p_i}_{v_i}(g) \text{ for all } f, g \in \{a, b, c\}^{\{\omega_1, \omega_2\}} \text{ such that } f \neq g\},$$

with $\mathcal{V}$ the set of normalized injective valuation functions v_i over the set X = {a, b, c} and $\mathcal{P}$ the set of possible injective probability measures over the set of possible events $2^\Omega = \{\emptyset, \{\omega_1\}, \{\omega_2\}, \{\omega_1, \omega_2\}\}$.
An SCF is a map $\varphi(v, p) : \mathcal{L}^{\{1,2\}} \to X^\Omega$.
If we hope to find some SCFs that do not suffer from the problems pointed at the end of the previous chapter, we have to relax the requirement of unanimity and introduce another axiom compatible with strategyproofness. While discussing the Top Selection Lemma in the previous section, we have seen that the introduction of unanimity excludes the possibility of any compromise results; only the top outcomes are selected. The following axiom imposes

that the SCF must select the compromise act (b, b) when the two agents have opposite preferences over outcomes and they both rank b as second.
We alleviate the notation by writing $v_i : a \succ b \succ c$ if $v_i(a) > v_i(b) > v_i(c)$, with $i = 1, 2$.

Compromise axiom.
For each $(v, p) \in \mathcal{L}^{\{1,2\}}$ *such that* $v_i : a \succ b \succ c$ *and* $v_j : c \succ b \succ a$, *with* $i \neq j$, *then* $\varphi(v, p) = (b, b)$.

To avoid repetitions, when a SCF respects this axiom we also say that the SCF *compromises*.
This axiom guarantees that at least in some of those situation in which a compromise act seems the best choice, it will indeed be selected. Notice that the axiom is pretty weak, since it only applies to a particular order over outcomes, namely when b is ranked second by both agents and they have opposite ordering over outcomes. Another important feature is that the axioms does not discriminate on agents; it does not depend on who have a particular preference order, as long as one prefers a to b to c and the other has the opposite ranking. Obviously, the axiom substitutes the problem of unanimity with another problem, namely that (b,b) will be selected at some profiles in which is would be more efficient to sect (a,c) or (c,a). This is the case when the order over outcomes for the two agents is as in the compromise axiom, but both agents think that b is a bad outcome. We will see in section 3.4 that this feature will have some implications both in terms of efficiency and for the requirement of group strategyproofness.
In addition, the axiom depends on the *order over outcomes*. In this way, this axiom will be compatible with the ordinal nature of strategyproofness in this domain (see the Ordinality lemma in section 2.2.1) and at the same time it will cover enough cases to allow for a characterization.
A different idea for could have been to make the profiles at which the SCF compromises depend on the *order over acts*. This could potentially reduce the inefficiencies introduced by the compromise axiom while still maintaining some normative appeal. With the following example we show that one particular type of such axioms is not compatible with strategyproofness. This example by itself does not show that there is no space for interesting results in this direction, but we think it gives the intuition that this idea may be less promising than the one we have adopted.

Example. The axiom that we are considering is the following:

Compromise axiom*. Take any profile at which the two agents have an opposite ordering over acts. At this profile the SCF must select the fifth-ranked act.

Since we have nine acts in this particular model, the fifth-ranked act will be the same for both agents when they have an opposite ordering over them. Furthermore, differently from the previous axiom, this one is *neutral over outcomes*, since it applies in all profiles at which agents have opposite preferences over acts, no matter the name of the outcomes.

Claim 3.1.1. *Compromise axiom* is incompatible with strategyproofness in this domain.*

Proof. To prove the claim we have to build a counterexample in which the application of strategyproofness and compromise axiom* leads to a contradiction.
Fix $p \in \mathcal{P}^{\{1,2\}}$. Consider some valuation functions with order over outcomes as shown in

following table. Assume that (v_1, p_1) and (v_2', p_2) generates opposite orders over acts and

φ	$v_1 : a \succ b \succ c$	$v_1' : c \succ a \succ b$
$v_2 : b \succ a \succ c$	$(b, b) \neq (a, a)!$	(a, a)
$v_2' : c \succ b \succ a$	(b, b)	

Table 3.1.1

that both agents rank the act (b, b) as fifth. Therefore, by compromise axiom*, the SCF must select (b, b) in this profile.
Assume also that (v_1', p_1) and (v_2', p_2) generates an opposite ordering over acts, with a the fifth outcome. By compromise axiom*, the social choice at this profile must be (a, a).
By monotonicity and permutation invariance, $\varphi((v_1, v_2'), p) = (b, b)$ implies that $\varphi((v_1, v_2), p) = (b, b)$. In the same way, by monotonicity and permutation invariance, if $\varphi((v_1', v_2), p) = (a, a)$ then $\varphi((v_1, v_2), p) = (a, a)$. By ordinality this leads to a contradiction, proving the claim.

□

3.2 Three lemmas on beliefs

The three lemmas on tastes presented in section 2.1.2 tell us what are the limitations imposed by strategyproofness to the social choice when there is a unilateral deviation of valuations. Nevertheless, agents can deviate in their beliefs, too. With the following three lemmas we derive further implications of strategyproofness when there is a unilateral change in beliefs. We call these *lemmas on beliefs* when we want to distinguish them from the previous three lemmas. Otherwise, we will refer to them as Lemma 1, Lemma 2 and Lemma 3. Until the end of this section φ denotes a strategyproof SCF.

We call $\varphi(v, p; \omega)$ the outcome selected by φ at state ω under the profile (v, p).

Lemma 1 on beliefs. *Consider $(v, p), (v, p') \in \mathcal{L}^{\{1,2\}}$ with $p = (p_i, p_j)$, $p' = (p_i', p_j)$ and $i \neq j \in \{1, 2\}$. If $\varphi(v, p) \neq \varphi(v, p')$, then*

$$\left[p_i'(\omega_1) > (<) p_i(\omega_1)\right] \Rightarrow \left[\begin{array}{ll} v_i(\varphi(v, p'; \omega_1)) & > (<) v_i(\varphi(v, p; \omega_1)) \textit{ and} \\ v_i(\varphi(v, p'; \omega_2)) & < (>) v_i(\varphi(v, p; \omega_2)) \end{array} \right].$$

This lemma is saying that if an agent increases his subjective probability that one state will be selected and this changes the social choice (everything else equal), then this this agent must receive a strictly better outcome in this state and a strictly worse outcome in the other state.

Lemma 2 on beliefs. *Consider $(v, p), (v, p') \in \mathcal{L}^{\{1,2\}}$ with $p = (p_i, p_j)$, $p' = (p_i', p_j)$ and $i \neq j \in \{1, 2\}$. If $\varphi(v, p) \neq \varphi(v, p')$, then either $p_i(\omega_1) > 0.5$ and $p'(\omega_1) < 0.5$ or $p_i(\omega_1) < 0.5$ and $p'(\omega_1) > 0.5$.*

Lemma 2 imposes that the only cases in which a unilateral change in beliefs changes the social choice are those in which an agent changes his beliefs on which of the two states of the world is most likely to happen.

Lemma 3 on beliefs. *Consider* $(v,p),(v,p') \in \mathcal{L}^{\{1,2\}}$ *with* $p = (p_i, p_j)$, $p' = (p'_i, p_j)$ *and* $i \neq j \in \{1,2\}$. *If* $\varphi(v,p) \neq \varphi(v,p')$, *then* $\varphi(v,p') = (\varphi(v,p;\omega_2), \varphi(v,p;\omega_1))$.

Namely if a unilateral change in beliefs changes the final outcome, at this new profile the social act must select at state ω_1 the outcome previously selected at state ω_2, and at state ω_2 the outcome previously selected at state ω_1.

The formal proofs of these lemmas can be found in Appendix A1.

3.3 The Characterization Theorem

The main result of our study is the following Theorem. This is a characterization of all the SCFs that respect strategyproofness and the compromise axiom.

Theorem 1. ***(Characterization)*** *An SCF* $\varphi : \mathcal{L}^{\{1,2\}} \to X^{\Omega}$ *is strategyproof and it compromises if and only if it is one of the following four kinds of functions:*

1. (b,b) *is selected at any profile of preferences;*

2. *if at least one agent prefers b to a(b to c), then the SCF selects* (b,b)*; when both agents prefer outcome a to b (c to b), then the SCF will select* (b,a) *or* (a,b) *(respectively* (b,c) *or* $(c.b)$*) in a way that can depend on agents beliefs (see Remark 3.3.1 for details);*

3. *if at least one agent prefers b to a(b to c), then the SCF selects* (b,b)*; when both agents prefer outcome a to b (c over b), then the SCF will select* (a,a) $((c,c))$*;*

4. *if both agents prefer a to b and at least one agent prefers b to c, then (b,a)((a,b)) is selected; if both agents prefer c to b and at least one agent prefers b to a then (c,b)((b,c)) is selected; if they both prefer a to b and c to b, then (c,a)((a,c)) is selected; otherwise (b,b) is selected.*

Notice that at each point we are using parenthesis to distinguish two different functions. For instance, at point 3 we describe the two functions: (3.a) when at least one agent prefers b to a, then the SCF selects (b,b) and when both agents prefer a to b the SCF selects (a,a); (3.b) when at least one agent prefers outcome b to c then the SCF selects (b,b) and when both agents prefer outcome c to b then the SCF selects act (c,c). We refer to function of type 3 or function 3 to refer to one of these two functions as we do for the other functions in Theorem 1.
To further grasp how these functions work, we have represented functions 1, 3 and 4 in Table 3.3.1. We did not represent function 2 because this also depends on the agents' beliefs in a way that we will explain after.

φ	v_1 a b c	v_1' a c b	v_1'' b a c	v_1''' b c a	$v_1'^v$ c a b	v_1^v c b a
$v_2 : a \succ b \succ c$	(b,b) (a,a) (b,a)	(b,b) (a,a) (b,a)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (a,a) (b,a)	(b,b) (b,b) (b,b)
$v_2' : a \succ c \succ b$	(b,b) (a,a) (b,a)	(b,b) (a,a) (c,a)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (a,a) (c,a)	(b,b) (b,b) (c,b)
$v_2'' : b \succ a \succ c$	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)
$v_2''' : b \succ c \succ a$	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)
$v_2'^v : c \succ a \succ b$	(b,b) (a,a) (b,a)	(b,b) (a,a) (c,a)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (a,a) (c,a)	(b,b) (b,b) (c,b)
$v_2^v : c \succ b \succ a$	(b,b) (b,b) (b,b)	(b,b) (b,b) (c,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (b,b)	(b,b) (b,b) (c,b)	(b,b) (b,b) (c,b)

Table 3.3.1

The red function corresponds to function 1 in the theorem; the blue function corresponds to a function of type 3; the green function corresponds to a function of type 4.
The following Remark explains how the functions of type 2 aggregate the information of agents' beliefs.

Remark 3.3.1. ***(On SCFs 2)***. These functions are the only strategyproof SCFs that compromise and that depend on beliefs in this restricted domain. To understand how beliefs can influence the social choice consider the following function.
Given $p \in \mathcal{P}^{\{1,2\}}$,

- if at least one agent prefers b to a, then the SCF will select (b, b);
- when both agents prefer outcome a to b, then the SCF will select (b, a).

First consider a $p' = (p_1', p_2)$ such that p_1' a passage from p to p' (*ceteris paribus*) changes the social choice. By Lemma 1 and Lemma 3 at those profiles in which the SCF selects (b, b), the outcome cannot change. Therefore, we have to look at those profiles in which both agents prefer a to b. By Lemma 3, the selection must pass from (b, a) to (a, b). This means means that the agents are receiving a strictly better outcome under state ω_1 and strictly worse outcome at ω_2. By Lemma 1 it must be that $p_1'(\omega_1) > p_1(\omega_1)$ and by Lemma 2 it must be that $p_1'(\omega_1) > 0.5$ and $p_1(\omega_1) < 0.5$.
By applying this reasoning to a deviation on the beliefs of agent 2 we can produce SCFs that depend on agents' beliefs and that do not contradict strategyproofness. We provide two example of such functions:

***Example 1*.** The social choice is decided by the following rule:

- if at least one agent prefers b to a, then the SCF will select (b, b);
- when both agents prefer outcome a to b *and* agent 1 prefers state ω_1 to state ω_2 then the SCF selects (a, b);
- otherwise (b, a) is selected.

***Example 2*.** The social choice is decided by the following rule:

- if at least one agent prefers b to a, then the SCF will select (b, b);
- when both agents prefer outcome a to b *and* both agents prefer state ω_1 to state ω_2 then the SCF selects (a, b);
- otherwise (b, a) is selected.

The proof of the characterization theorem (see Appendix A.2) is made of two main steps. We first prove that strategyproofness and the compromise axiom imply that the SCF must be one of those in Theorem 1 (proof of sufficiency); only then we show that all those functions respect strategyproofness and the compromise axiom (proof of necessity).
For the first part of the proof we fix the beliefs to a profile $p \in \mathcal{P}^{\{1,2\}}$. This assumption will be relaxed at the end of the proof. Then, we derive Lemma 4 and Lemma 5.

Lemma 4. *When b is the top alternative of at least one agent, then φ must select (b, b).*

Lemma 5. *For each $v = (v_1, v_2) \in \mathcal{V}^{\{1,2\}}$ and $w = (w_1, w_2)$ with $v_1 = w_2$ and $v_2 = w_1$, then $\varphi(v, p) = \varphi(w, p)$.*

Lemma 5 states that, fixed a profile of beliefs $p \in \mathcal{P}^{\{1,2\}}$, then the social choice at any profile (v, p), for any $v \in \mathcal{V}^{\{1,2\}}$, should not depend on the name of the agents. In other words, the SCF must be anonymous with respect to valuations.
The proof is than developed by exhaustion, considering all the possible selections at two particular profiles of preferences. We than relax the assumption about the beliefs and study, by using the three lemmas on beliefs, how a change in belief can impact the social choice for each one of the previous cases. This will lead to identify the four kinds of SCFs of Theorem 1.
In the second step of the proof, i.e. showing that these functions respect strategyproofness and the compromise axiom, we simply proceed by cases.

3.4 Comparison between the compromise result and BS

In this section we compare the *locally bilateral top selection* SCFs of BS with the functions that we have characterized in Theorem 1 on three different dimensions. We first ask whether these two classes of functions can be seen as different ways of selecting social acts. To stress the difference between these two functions, we consider different properties that distinguish them. Second, we compare our function 4 with a *bilaterally consensual* rule in terms of efficiency. We end the section by pointing to some implications of the stronger requirement

of *group-strategyproofness* for these two classes of functions.
That our result constitute an intrinsically different way of taking group decisions in context of uncertainty can be made clear by points out some of the features in which these functions differ. First and most evidently, they rely on different kind of information about agents' preferences to produce a social decision. On the one hand, we have seen that the Top Selection Lemma in BS limits the relevant information concerning the agents' preferences over the outcomes to their top alternative; all the other information about the outcomes is discarded. On the contrary, our functions are more sensible to the agents' preferences over outcomes. For instance, to make a decision under Function 4 we need to know how the agents rank two out of the four couple of outcomes. On the other hand, our functions, with the exception of function 3, do not consider the agents' beliefs, whereas the selection rules in BS rely heavily on this kind of information. This dichotomy may also suggest that the functions that we have characterized may cover a big portion of the strategyproof functions that are left behind by unanimity, at least this in this restricted model. In other words, it may suggest that the compromise axiom is complementary to unanimity.
Moreover, by assumption these two classes of functions differ in that ours compromise and the functions in BS do not (as explained in the 2.1.2). However, this is not the only interesting normative feature that distinguishes the two class of functions. As already stressed in BS, their functions are not *anonymous*. As long as two agents have different top outcomes, if we change the name of two agents, the act selected by any locally bilateral top selection rule changes. In our case, instead, almost all of our functions are anonymous. The only non anonymous functions are those function of type 3 that make the choice contingent to the belief of one particular agent, discarding the other agent's beliefs (for an example see Example 2 in the previous section). Since the model in this paper finds most of its applications in situations of voting, anonymity seems a natural requirement to meet. The identification of anonymous and strategyproof functions in this domain was already been notices by BS as a important problem left open by their analysis. We have provided an answer, even though a limited one.
It is also interesting to notice that the functions in BS are *neutral* among outcomes, while our functions are not. Our compromise axiom is skewed towards outcome b, that is selected more often then the other outcomes, while in the case of BS, their functions do not depend in any respect on the name of the outcome. We consider this a less important normative requirement to meet compared to anonymity, but it is still relevant that these functions show a dichotomy even in this respect.
The other question that we want to analyse in this section is how much efficiency we have to renounce to have a strategyproof functions that compromises compared to a function that does not. In particular, we compare our most efficient function, namely function 4, with the most efficient rule in BS, that is the bilaterally consensual rule.
Before proceeding with the analysis, we have to recognize that in any case with two states of the world, there is no composite locally bilateral top selection, but there are only three kind of SCFs that correspond to the three kinds of selection rules. This is because there are only two possible partitions of the state space, namely $\Omega' = \{\{\omega_1\}, \{\omega_2\}\}$ and $\Omega'' = \{\{\omega_1, \omega_2\}\}$. If we associate a bilaterally dictatorial or consensual rule to a set that contains only one state, then these rules will coincide with the constant assignment. Therefore the SCF that we consider here consider the partition Ω'' and assigns the two states by a bilateral consensual rule (for more details on how this rule works see section 2.1.2).
We undertake this analysis by considering the following three cases:

(a) the two agents agree on their top outcome, that is different from b;

(b) the two agents disagree on their first ranked outcome and one of them ranks b first;

(c) the two agents disagree on their first ranked outcome and no one ranks b first.

CASE (a)

Without loss of generality, consider the case where both agents rank a first. The following table represents all the possible profiles of preferences of this kind together with the social acts selected by the bilaterally consensual rule and by our function 4.

φ	I		II		III		IV	
	v_1	v_2	v'_1	v'_2	v''_1	v''_2	v'''_1	v'''_2
	a	a	a	a	a	a	a	a
	b	b	c	c	c	b	b	c
	c	c	b	b	b	c	c	b
bilaterally consensual	(a, a)		(a, a)		(a, a)		(a, a)	
function 4	(b, a)		(c, a)		(b, a)		(b, a)	

Table 3.4.1

In all these situations there is a loss in efficiency for all the agents. This is due to the the fact that our rules are not unanimous. Nevertheless, this loss may not be too large. The preferred outcome is still selected in one of the two states of nature and the magnitude of the loss depends on the actual value that both agents associate to the outcome selected at the other state.

CASE (b)

Without loss of generality consider a bilaterally consensual rule that provisionally associate state ω_1 with agent 1 and state ω_2 with agent 2. Assume also that agent 2 ranks a first. The possible social acts are represented in Table 3.4.2.

φ	$p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$	Otherwise
bilaterally consensual	(a, b)	(b, a)
function 4	(b, b)	(b, b)

Table 3.4.2

We can see that function 4 compared to the other rule welfare improves the situation of agent 1 in both cases, while it makes agent 2 worse off.
We can calculate this welfare improvement for both agents. When $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$ the the welfare gain for agent 1 is given by

$$p_1(\omega_1)(1 - v_1(a)).$$

In the other profiles, this gain is given by

$$p_1(\omega_2)(1 - v_1(a)).$$

For player 2, when $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$ the above welfare loss is

$$p_2(\omega_1)(v_2(b) - 1).$$

In the other profiles, this loss is

$$p_2(\omega_2)(v_2(b) - 1).$$

By adding the welfare gain of agent 1 to the welfare loss of agent 2 we can calculate the welfare difference between function 4 and the bilaterally consensual rule in all these profiles. When $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$, then this welfare difference is given by

$$p_1(\omega_1)(1 - v_1(a)) + p_2(\omega_1)(v_2(b) - 1).$$

In all the other profiles it is given by

$$p_1(\omega_2)(1 - v_1(a)) + p_2(\omega_2)(v_2(b) - 1).$$

A low enough value of $v_1(a)$ and a high enough value of $v_2(b)$ would produce a welfare gain for function 4 compared to the bilaterally consensual rule. These are situation of compromise in which both players agree that one outcome (outcome b) is good, but they disagree on another outcome (outcome a). It is therefore not surprising that our functions do well in terms of efficiency in these situations.

CASE (c)

Consider without loss of generality the case in which agent 1 ranks a first and agent 2 ranks c first. Assume also that agent 1 is provisionally assigned to state ω_1 and agent 2 to state ω_2 by the bilaterally consensual rule. Table 3.4.3 shows the possible profiles of valuations of this kind together with the social choices for the two different functions. For the bilaterally consensual rule, the left social choice corresponds to the case $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$, the right to the other profiles of beliefs.

	I		II		III		IV	
	v_1	v_2	v_1'	v_2'	v_1''	v_2''	v_1'''	v_2'''
φ	a	c	a	c	a	c	a	c
	b	b	c	a	b	a	c	b
	c	a	b	b	c	b	b	a
bilaterally consensual	(c,a)	(a,c)	(c,a)	(a,c)	(c,a)	(a,c)	(c,a)	(a,c)
function 4	(b,b)		(c,a)		(b,a)		(b,a)	

Table 3.4.3

We study the welfare difference between function 4 and the bilaterally consensual rule as we did in the previous case for each one of the four profiles of valuations.

Profile I. When $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$, then the welfare difference is

$$(v_1(b) + v_2(b)) - (p_1(\omega_2) + p_2(\omega_1)).$$

For the other profiles of beliefs, then the welfare difference is

$$(v_1(b) + v_2(b)) - (p_1(\omega_1) + p_2(\omega_2)).$$

Notice that at Profile I the compromise axiom directly applies. Therefore, if both agents value the middle alternative high enough, then our function can produce a welfare gain on the bilaterally consensual rule.

Profile II. When $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$ the social choice remains the same for both functions. When $p_1(\omega_1) > 1/2$ and $p_2(\omega_1) < 1/2$, then the bilaterally consensual rule welfare improves on function 4. In the other two cases remaining there is one agents that is better off under function 4 and another one that is worse off. We can see here that our function 4 does as well as random dictatorship in terms of efficiency.

Profiles III. When $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$, then agent 1 is strictly better off under function 4, while agent 2 is strictly worse off. The agents' beliefs and agent 1's valuation of outcome b determines whether function 4 will produce welfare improvement or loss on the bilaterally consensual rule.
In the other cases the welfare difference is given by

$$p_1(\omega_1)(v_1''(b) - 1) + p_1(\omega_2) + v_2''(a)(p_2(\omega_2) - p_2(\omega_1)) - p_2(\omega_2).$$

Even in this case, we can see that function 4 does not necessarily lead to a welfare loss compared to the bilaterally consensual rule.

Profiles IV. Under $p_1(\omega_1) < 1/2$ and $p_2(\omega_1) > 1/2$, function 4 makes both agents worse off compared to the other function. The size of the loss is determined by the value that agent 1 gives to outcome c and by the value that agent 2 gives to outcome b.
In the other profiles of beliefs the welfare difference between function 4 and the bilaterally consensual function is given by

$$p_1(\omega_2)(1 - v_1'''(c)) - p_1(\omega_1) + p_2(\omega_1)v_2'''(b) - p_2(\omega_2).$$

As before what will determine whether there will be a welfare loss or not is the actual values attributes by the two players to the outcomes and their beliefs.

This analysis makes clear that, as far as efficiency is concerned, the key difference between the two functions is their responsiveness to the agents' preferences over outcomes. We know by the ordinality lemma that if a strategyproof function wants to aggregate this information in a way that takes into account more than the agents' top alternative, then it must do it by only considering the preferences order. Function 4 is more responsive than the bilaterally consensual rule to agents' preferences over outcomes and this inevitably introduces inefficiencies. In the particular case of the bilaterally consensual rule, there are also some efficiency gains from the way this function aggregates beliefs. This is particularly evident in Profile II of case (c) in which the possibility of switching the assigned state between the agents allows the bilaterally consensual function to catch up some inefficiencies on function 4.
The welfare loss of function 4 are reduced or reversed in those situations in which the two agents would prefer to receive a compromise act rather than the other agent's top alternative in one state of the world. This is a consequence of the compromise axiom that also holds in those profiles in which the axiom does not directly apply, e.g. case (b). Overall, we think that the loss in efficiency of function 4 compared to the bilaterally consensual rule are not too high to be of much concern.

We want to close this section by considering a stronger requirement of non manipulability, namely *group strategyproofness.*
An SCF is *group strategyproof* if for all $S \subseteq N$, $i \in N$, $(v,p) \in \mathcal{D}^N$, $v'_S \in \mathcal{V}^S$ and $p'_S \in \mathcal{P}^S$,

$$E^{p_i}_{v_i}(\varphi(v,p)) \geq E^{p_i}_{v_i}(\varphi((v'_S, v_{N \setminus S}), (p'_S, p_{N \setminus S}))).$$

Group strategyproofness requires that there must be no group of agents that can all strictly benefit from a cooperative manipulations of preferences. This requirement is obviously stronger than that of individual strategyproofness that we have used throughout this study. However, in many situations, requiring individual strategyproofness is not enough if there is a group of people that can manipulate the social choice in their favour. Hence, it is natural to ask whether there are some functions in the two classes that we have studied that also respect this stronger property of non manipulability.
Concerning the class of *locally bilateral top selection* SCFs, it is possible to show that the only group strategyproof SCF is the dictatorial one, i.e. the function that assigns all states of nature to one agent. That there is no group strategyproof SCFs composed of bilaterally dictatorial or bilaterally consensual rules comes from the fact that these functions do not compromise. In fact, it is always possible to find a profile of preference in which the agents would prefer the compromise act to the one that assigns their top outcome to some states of nature; this leave open the opportunity for a beneficial cooperative manipulation of preferences. Moreover, the constant assignment rules can be manipulated by a group of agents when these are assigned to states of nature that they would prefer to exchange with each other.[1]
Since the previous problem of compromise is solved by our function, we could expect to find some nontrivial group strategyproof functions in the class characterized by Theorem 1. In fact, it can be easily proven that functions 1, 2 and 3 in Theorem 1 are group strategyproof and that function 4 is not. This last point is a direct consequence of our compromise axiom. To see this consider the function of type 4 that selects (c,a) when both agents prefer c and a to b. Consider a situation in which the compromise axiom applies, for instance $v_1 : a \succ b \succ c$ and $v_2 : c \succ b \succ a$. At this profile, no matter the beliefs and the actual valuations, the social choice is (b,b). Nevertheless, if $v_i(b)$ for $i = 1,2$ is low enough and $p_1(\omega_1) > 0.5$ and $p_2(\omega_1) < 0.5$, than both players would prefer (c,a) to (b,b). They can get if both of them appropriately misreport their preferences.
We can conclude that group strategyproofness is a very restrictive requirement in this domain of preferences. Even though our functions 1, 2 and 3 respect this requirement, we must recognize that these are not the most interesting functions, among those that we characterized. Functions of type 4 have more interesting properties that the previous three, but they are not group strategyproof.[2]

[1] Proofs can be provided upon request.

[2] Another way to study group strategyproofness is to look at the domain of preferences. The important paper Barberà, Berga, and Moreno 2010 gives the sufficient conditions that a domain of preferences must respect for individual strategyproofness to imply group strategyproofness in the case of pure public goods. They show that these conditions become also necessary if the domain of preferences allow for opposite preferences. It can be shown that SEU preferences, because they are separable, always allow for opposite preferences. Therefore, as implied by the previous arguments, the above conditions on the domain of preferences are not respected. Proofs can be provided upon request.

Chapter 4

Conclusion

In this study we have characterized a new class of strategyproof SCFs in the domain of SEU preferences. To do so required to developed further the study of the restrictions that strategyproofness imposes on this domain of preferences. This led us to produce the three lemmas on beliefs, that were then used in the proof of Theorem 1. The class of functions that this theorem characterizes shows different properties compared to the *locally bilateral top selection rules* in BS and can be seen as an alternative rule to make group choices in situations of uncertainty.

In particular, our functions are able to extrapolate more information about agents' preferences over outcomes. For instance, our function 4 aggregates the information about agents' preferences over two of the three possible couples of outcomes. Moreover, most of the functions in Theorem 1 are anonymous, which we see as an important normative requirement in any situation of voting. Finally, as we have shown in section 3.4, they do not imply a big loss in efficiency compared to the functions in BS.

We want to spend the space that remains in trying to partially address a question that naturally emerges from our analysis, namely the possibility of extending the analysis to a more general model.

Enlarging our model with more states of nature does not create many problems. The compromise axiom can still be used and much of the reasoning in the proof of Theorem 1 still applies. This comes from the fact that for most of the proof we have fixed the subjective beliefs to a profile $p \in \mathcal{P}^{\{1,2\}}$. For instance, Lemma 4 and Lemma 5 immediately apply to cases with more states of nature. Obviously, if we want to extend this model to any number of states of nature, we cannot anymore proceed by exhaustion and we need to generalize the proof of Theorem 1. We claim that with the right choice of definitions and notations, this is a manageable task.

Nevertheless, once we add more states of nature to the model, the three lemmas on beliefs do not immediately hold in the way we presented them here. Similar results could potentially be derived adequately substituting the two states ω_1 and ω_2 with an event E and its complementary event E^c. This results would be more general but less specific than the previous ones. Therefore, they must be used inside a more general argument than the one in this study.

If, in addition, we want to extend the model to any number of outcomes and to any number of agents, then we need a different compromise axiom. We think that the new axiom should have three features in common with the one that we have used. First, it must be ordinal. As we have shown in section 3.2 (see compromise axiom*), using a non ordinal axiom is at odd with the ordinal nature of strategyproofness. Second, it must impose a compromise act under certain profiles of preferences. Third, it must compromise only in relation to a particular outcome; in our case this was outcome b. This last feature comes from the realization that a compromise axiom that is neutral among outcomes may be too restrictive. The following is a suggestion for such an axiom.

Call v_i^j the value that agent i attributes to agent j's top outcome and write $\bar{b} = (b, ..., b)$, the act that that selects outcome b in any state of nature. Consider a set N of agents.

Compromise axiom**. If $v \in \mathcal{V}^N$ is such that $v_i(b) > v_i^j$ for all $i \neq j \in N$, then $\varphi(v, p) = \bar{b}$ for all $p \in \mathcal{P}^N$.

The previous axiom is saying that whenever all agents prefer outcome b to any other agent's top alternative, then the SCF must select b in all states of nature.
It can be shown that if we fix the number of agents to two and we increase the number of outcomes, by applying the previous axiom and strategyproofness, the only profiles of preferences that we have to study are of three kinds: those in which both agents agree that b is their worst outcome; those in which one agent ranks b second-last and the other ranks it last; those in which both agents rank b second-last and they disagree on their worst outcome. In all other profiles of preferences, the social choice function must select $\bar{b}$.[1] Then, even though we have to decide among more profiles of preferences than before, an increase in the number of outcomes does not make the problem intractable and a generalization of our reasoning should be able to approach a solution.
Finally, as the literature on strategyproofness suggests (see Sen 2001), a solution for the case with two agents should not lie too far from a solution for the general case of any number of agents.

[1]This implication comes from arguments similar to those applied in the proof of Theorem 1. A proof can be provided upon request.

Bibliography

Arrow, K.J. (1971). *Essays in the Theory of Risk-bearing*. Markham economics series. North-Holland.

Bahel, Eric and Yves Sprumont (2020). “Strategyproof Choice of Social Acts”. In: *American Economic Review* 110.2, pp. 596–627.

Barberà, Salvador (2011). “Strategyproof social choice”. In: *Handbook of social choice and welfare*. Vol. 2. Elsevier, pp. 731–831.

Barberà, Salvador, Dolors Berga, and Bernardo Moreno (2010). “Individual versus group strategy-proofness: When do they coincide?” In: *Journal of Economic Theory* 145.5, pp. 1648–1674.

Drèze, J. (1990). *Essays on Economic Decisions Under Uncertainty*. Cambridge University Press.

Drèze, Jacques H and Aldo Rustichini (2004). “State-dependent utility and decision theory”. In: *Handbook of utility theory*. Springer, pp. 839–892.

Gibbard, Allan (1973). “Manipulation of voting schemes: a general result”. In: *Econometrica: journal of the Econometric Society*, pp. 587–601.

— (1977). “Manipulation of Schemes that Mix Voting with Chance”. In: *Econometrica* 45.3, pp. 665–681.

Haller, Hans (1985). “Expected utility and revelation of subjective probabilities”. In: *Economics Letters* 17.4, pp. 305–309.

Hylland, Aanund (1980). “Strategy proofness of voting procedures with lotteries as outcomes and infinite sets of strategies”. In: Unpublished paper, University of Oslo.[341, 349].

Mongin, Philippe (2016). “Spurious unanimity and the Pareto principle”. In: *Economics & Philosophy* 32.3, pp. 511–532.

Moulin, H. (1980). “On Strategy-Proofness and Single Peakedness”. In: *Public Choice* 35.4, pp. 437–455.

Satterthwaite, Mark A. (1975). “The existence of strategy proof voting procedures”. In: Ph.D dissertation. University of Winsconsin-Madison.

Savage, LJ (1954). *The Foundations of Statistics*. New York, NY, John Wiley & Sons Incorporated.

Sen, Arunava (2001). “Another direct proof of the Gibbard–Satterthwaite theorem”. In: *Economics Letters* 70.3, pp. 381–385.

Sprumont, Yves (1995). “Strategyproof collective choice in economic and political environments”. In: *Canadian Journal of Economics*, pp. 68–107.

Zeckhauser, Richard (1973). “Voting systems, honest preferences and Pareto optimality”. In: *The American Political Science Review* 67.3, pp. 934–946.