

Social Choice with Uncertain Preferences

Bruno Strulovici*

Nuffield College, Oxford University

April 3, 2007

Extended Abstract

The social choice literature either assumes that individuals know their preferences or, at the extreme opposite, that they are placed behind a veil of ignorance. In reality, preferences are to some degree uncertain and may be learned through experimentation. For example, a reform may have uncertain consequences which can only be learned by implementing the reform to some extent. Similarly, decisions in committees are often analyzed under the assumption that payoff distributions are known to committee members at the time of their decisions, although this is rarely the case in practice.

This paper considers how uncertain preferences may affect collective decisions. The question is analyzed in the context of a multi-armed bandit model, each arm of which may have contrasted effects across society members. At each instant, society chooses a single action according to some fixed voting rule. In addition to a “safe” status quo action, society can experiment with riskier actions, which allows a better assessment of individual valuations for such actions, hence a better knowledge of one’s preferences.

Under any voting rule and with independently distributed preferences, risky actions are shown to be less valuable than in the case of a single decision maker, due to a *control-loss effect*. As a result, society experiments less than each individual would in a single-decision making problem with otherwise identical conditions. Even if everyone ignores this control-loss effect, experimentation under a majority rule is always lower than what would be optimal for a utilitarian society, as the option value of risky actions creates an *experimentation bias* in their favor, which is not captured by the majority rule.

*Correspondance should be sent to Bruno Strulovici, Nuffield College, Oxford OX1 1NF, U.K. Email: bruno.strulovici@economics.ox.ac.uk. Phone +44(0)1865-278639.

In a specialized, two-arm setting, the severity of this discrepancy diminishes if preferences are more likely to be similar. Preference similarity may also make experimentation more attractive if learning is accelerated by individuals observing both their and other individuals' payoffs.

When individuals are *ex ante* identical and positive information about the risky action arrives at discrete times, a dynamic Rawls criterion emerges endogenously under majority rule, whereby experimentation stops when those that value the risky action the least become indifferent between that action and the status quo. However, if a decisive number of individuals learn enough positive news before the Rawls level of experimentation is reached, the risky action is imposed forever to society. With such information jumps, the Rawls criterion is thus violated *ex post* when the risky action is imposed to individuals who voluntarily extended experimentation but turned out to be unlucky. The ability to pre-commit to a stopping rule such as the *ex post* Rawls criterion, which is not dynamically enforceable by a majority rule, may allow for more efficient levels of experimentation. These results thus embed the normative justification of Rawls criterion into a positive, dynamic analysis.

With multiple risky alternatives, correlation across alternatives can be exploited to conduct gradual reforms from safer to riskier actions. In general however, Condorcet cycles may emerge as individuals acquire precise information about their preferences.