CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1560

The Roman Metro Problem: Dynamic Voting and the Limited Power of Commitment

Christian Roessler* Sandro Shelegia* Bruno Strulovici**

March 2, 2013

JEL Classification: D70, H41, C70

*University of Vienna
**Northwestern University

The Roman Metro Problem: Dynamic Voting and the Limited Power of Commitment*

Christian Roessler[†] Sandro Shelegia[†] Bruno Strulovici[‡]

March 2, 2013

Abstract

A frequently heard explanation for the underdeveloped metro system in Rome is the following one: "If we tried to build a new metro line, it would probably be stopped by archeological finds that are too valuable to destroy, so the investment would be wasted." This statement, which seems self-contradictory from the perspective of a single decision maker, can be rationalized in a voting model with diverse constituents. One would think that commitment to finishing the metro line (no matter what is discovered in the process) can resolve this inefficiency. We show, however, that a Condorcet cycle occurs among the plans of action one could feasibly commit to, precisely when the metro project is defeated in step-by-step voting (that is, when commitment is needed). More generally, we prove a theorem for binary-choice trees and arbitrary learning, establishing that no plan of action which is majority-preferred to the equilibrium play without commitment can be a Condorcet winner among all possible plans. Hence, surprisingly, commitment has no power in a large class of voting problems.

JEL: D70, H41, C70

1 Introduction

The Roman underground has two lines, 49 stops, that serve a metropolitan area of 3.4 million residents, and 9 million annual visitors. Berlin is similar in size, but has 173 subway stations and an extensive overground train system. In fact, there is no major metropolitan area in Europe with more inhabitants per metro stop than Rome.[1] What is lacking underground cannot be compensated on the surface - the eternal city was not built with suspended monorails and large buses in mind. So why don't Romans invest in expanding their metro system? Every Roman will eagerly explain why: metro lines are extremely expensive, and in Rome, more than anywhere else in the world, one is likely to

*We thank Jeff Ely, Michael Greinecker, Karl Schlag, Stephen Schmidt, the editor, three anonymous referees, and seminar audiences for their comments.

[†]University of Vienna

[‡]Northwestern University

[1]Athens, whose population of 3.3 million is served by 33 subway stations, also has the historic Athens-Piraeus Electric Railway, which adds 24 more stops.

run into ancient ruins of such value that the metro project would be suspended to preserve the ruins.[2] Hence, Romans stopped trying to build metro lines altogether.

There is a problem with this explanation, illustrated by the following analogy. You would like to dig a well in a place where gold may be hidden. It is possible that you will strike upon gold while excavating, get rich, and never finish the well. Given this "risk," you decide not to dig in the first place. It seems like a foolish conclusion because the option value of finding gold is positive (you can, after all, go right past the gold if you want to). With the possibility of discovering gold, digging the well is more, not less, attractive. From this standpoint, Romans should not worry that they might discover ancient ruins. Ruins will only be preserved if they are more valuable than the metro line - and the construction project has paid off more than initially expected. Only a time-inconsistent person would refuse to start a project for fear of not finishing it because an even better opportunity might appear in the process. However, such time inconsistency can easily arise as a result of collective decision making through majority voting, as we demonstrate in the first part of the paper.

It is natural to ask whether commitment to future actions could fix the problem: if completion of the metro line could be guaranteed at the outset (regardless of whether ruins are found), would a majority then support the project?[3] The answer is, quite strikingly, negative. Instead, a Condorcet cycle arises when all feasible ways to commit are put to the vote. This observation in the context of the metro game is just an instance of a deep relationship between pathological voting outcomes and Condorcet cycles. The paper's main insight is an equivalence theorem that applies to dynamic majority voting over a finite time horizon (allowing for taste shocks and learning, which affect voter preferences in heterogeneous ways). In essence, our result says that whenever the equilibrium outcome can be improved by commitment, majority voting over all feasible ways to commit is indecisive: there is no Condorcet winner.

The dynamic voting literature emphasizes that voters will make today's choice based on how it will influence future voting outcomes. They must, in Penn's [10] words, be "farsighted." For example, in deciding whether to grant voting rights to additional persons, current voters will take into account that the new voters would create new majorities (Jack and Lagunoff [6]). Barbera et al. [1] showed that such considerations can lead to counterintuitive behavior, such as enfranchising one's enemies, rather than one's friends. Similarly, a majority of voters may oppose socially desirable projects, such as a reform in Fernandez and Rodrik [5] and experimentation in Strulovici [15] that offer net benefits to the population. Our stylization of the Roman metro problem adds to these another stark example of how voters may fail to ratify a socially desirable project, paradoxically because of the prospect of gaining a socially valuable option.[4] In a further application we present

[2]From an engineering standpoint, the issue is actually not the underground tunnels themselves, which could be built below the archeological layer; it's the various vents and exits that have to be dug at regular distances.

[3]In 2007, then-Chairman of Roma Metropolitane, Enrico Testa, has complained "there are treasures that are underground that would stay buried forever, but as soon as we uncover them, our work gets blocked." (In the Wall Street Journal, January 27, 2007, sec. A, p. 1.)

[4]Messner and Polborn [8] also discuss how an option (to delay) may lead to an inferior collective choice. Strulovici's ([14], see Section 7) story of three friends who vote (not to) sample an ethnic restaurant,

in the paper, inspired by the academic job market, voters even reject a Pareto-improving series of actions that is available to them.

Hence, dynamic voting often leads to an outcome that is, for a majority of the population, suboptimal compared to the outcome or distribution over outcomes that could have been attained through another set of choices. A subset of the population is deterred from voting along its preferred path by what may be decided at the next stage, that is, by lack of commitment. Our paper is the first to make the connection from these situations, where commitment is required to improve on a bad outcome, to a Condorcet cycle among all feasible ways to commit. This relationship is important because, if we allow commitment, then we should allow commitment to any plan of action that can be formulated (although we discuss a caveat later in the context of Fernandez and Rodrik's application to reform).

Although Condorcet cycles have previously been studied in dynamic settings (see Roberts [12], and Bernheim and Slavov [2]), we start with a dynamic voting game where majority preference is in general not cyclical (neither in direct choices, which are limited to two alternatives, nor over the terminal states). We then consider a static voting game that is derived from the dynamic game by pitting all possible sequences of actions against each other. The existence of cycles in the static game depends on a property of the equilibrium in the dynamic game.

Certain static voting "anomalies" are known to be linked to Condorcet cycles. The Ostrogorski (and closely related Anscombe) paradox notes that issue-by-issue comparison of candidates (the winner being the candidate who wins on most issues) can lead to a different ranking than voting directly on candidates (where individual voters rank them based on issues). Rae and Daudt [11], Kelly [7], and Nurmi and Saari [9] have shown that, with at least three candidates, the presence of the Ostrogorski paradox implies that there is a Condorcet cycle in pairwise voting. Our result also relates paradoxical voting outcomes to Condorcet cycles, but it is not driven by aggregation. The key is the expansion of the choice set in transitioning from a dynamic to a static setting. In this sense, it is kindred to Zeckhauser [16] (see also the comment by Shepsle [13]), who showed that, in general, when a majority-ranked set of alternatives is completed by including all lotteries over the set, no Condorcet winner exists. Considering plans of action in a dynamic setting can be viewed as another "natural" expansion that leads, under certain circumstances, to a Condorcet cycle.

We begin (in Section 2) with a discussion of the Roman metro problem as an example of how dynamic voting can "go wrong" in the absence of commitment.[5] If an initial majority decision is taken to build a metro line, an antiquity will be found with some probability, which triggers a second vote to determine whether construction should continue. The three possible outcomes are ranked by majority preference: status quo third, the metro line second, and the antiquity first. Moreover, utilitarian payoffs are aligned with majority

because the preferences of a non-pivotal member of the group would change so as to make him pivotal, shares the main features of our examples. We reproduce this example in Appendix B.

[5] While construction would in reality not be stopped by popular vote - rather, Rome's archeological authority, the Sopraintendenza ai Beni Culturali (Superintendency of Cultural Heritage) has the prerogative - one might argue that the existence of such institutions ultimately has to reflect majority opinion.

preference: status quo yields less total utility to the population than the metro line,[6] which yields less than the antiquity. Therefore, a social planner would be in favor of the metro project, the more so as finding ruins would create a valuable option.

Nevertheless, it turns out that a majority of voters may oppose the project. Some voters reject it outright because they find the metro line too costly, and would favor stopping if an antiquity were found. Metro proponents anticipate that some who are initially in their camp would join these metro opponents in voting to preserve the antiquity. If this outcome is sufficiently likely, even some who benefit from the metro vote against a project that may never produce the line.[7] We go on to show that, even if voters could commit to future actions, it is not possible to agree on a plan to complete the metro line. Allowing for a much broader class of (finite) voting games (in Section 3), we then derive the fundamental result on Condorcet cycles that arise with commitment (in Section 4), and illustrate it in some other dynamic voting games (in Section 5). We also discuss extension to infinite-horizon settings (in Sections 6 and Appendix B).

2 The Roman metro game

The city has the option to build a metro line, but there is a probability q that, while doing so, an antiquity will be found that lies in the path. If so, the antiquity could either be destroyed, or the construction project could be abandoned. The decision maker (mayor) is acting in line with majority preferences. Part of the consideration in building the metro is that it will have to be financed (through taxes or cuts elsewhere). Therefore, negative valuations are entirely possible. Suppose that the population consists of three types (A, B, C), who are present in equal numbers. A wants to build the metro, but only if it will not be abandoned if an antiquity is found. B wants to build the metro and abandon it if an antiquity is found. C does not want to build the metro, but if it is built and an antiquity is found, wants to abandon it. Thus, we have the following game tree, with some possible payoffs that reflect the preferences of the three constituencies:

Figure 1 goes here.

Note that, if $q = 0$ (that is, there is no possibility of finding an antiquity), building the metro has majority support (from A and B) and yields a surplus of 1 over the status quo. If an antiquity is found, a majority (B and C) will vote in favor of abandoning the metro.[8] Abandoning the metro not only has majority support, but also is an option that is strictly valuable in terms of total utility (yielding 2 instead of 1). Anticipating future voting outcomes, A initially votes for the project if her expected utility from the project is non-negative

$$E(u_A) = -2q + 1 - q \geq 0 \iff q \leq \frac{1}{3}.$$

[6]By total utility we mean an unweighted sum (integral) of utilities over the whole population.

[7]In the appendix, we provide more general conditions on the type distribution that cause the project to be defeated in the metro game.

[8]Formally, as the solution concept we use voting equilibrium in weakly undominated strategies. This amounts to solving the game with backward induction. For the formal definition see Section 3.

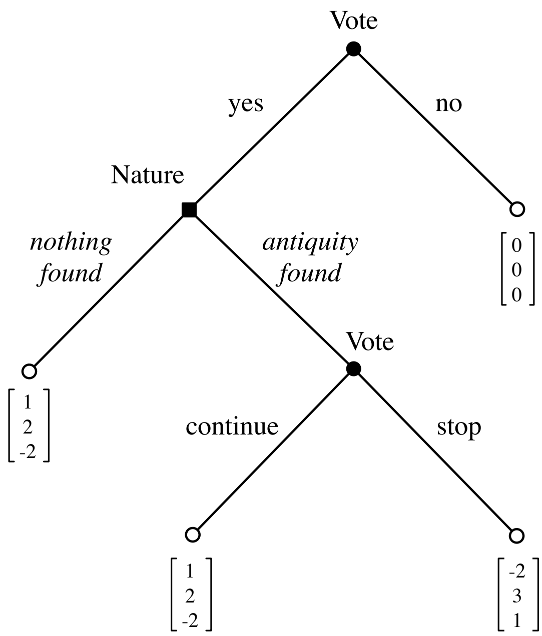

Figure 1: The Roman Metro Game. At the circular nodes, decisions are made according to voting majorities. The square node is a chance node.

B votes for the project regardless of q, because B benefits whether or not an antiquity is found. C votes for the project if

$$E(u_C) = q - 2\,(1 - q) \geq 0 \iff q \geq \frac{2}{3}.$$

Overall, there is a majority for the project at the outset if $q \leq 1/3$ (in which case, it is supported by A and B) or $q \geq 2/3$ (then, the project is supported by B and C). But in case $1/3 < q < 2/3$, a majority (consisting of A and C) will oppose it (see Figure 2 for illustration).

Figure 2 goes here.

At the same time, for any q, the total expected utility from the project is $2q + 1 - q = 1 + q$. This is always positive, so a utilitarian social planner would start the project regardless of q, and stop if an antiquity were found. Hence, adding the option to abandon may lead to an economically inefficient decision if q is intermediate. The problem is that the probability of finding an antiquity in Rome is relatively high, yet not high enough for anyone so inclined to support the project only as a way to search for antiquities.

The three-types example is not particularly special in generating a vote against the metro. We can replicate this choice in a voting game where the sequence of decisions is as above, but the population may be of arbitrary size, and payoffs satisfy restrictions that stack the deck squarely in favor of construction. We use M (metro) to denote the outcome that the metro is completed,[9] and T (antiquity) for the outcome that an antiquity is preserved (with the consequence that the metro line is never finished) and S for the

[9]For simplicity, we do not distinguish between the outcome metro after antiquity was found and destroyed and metro after no antiquity was found.

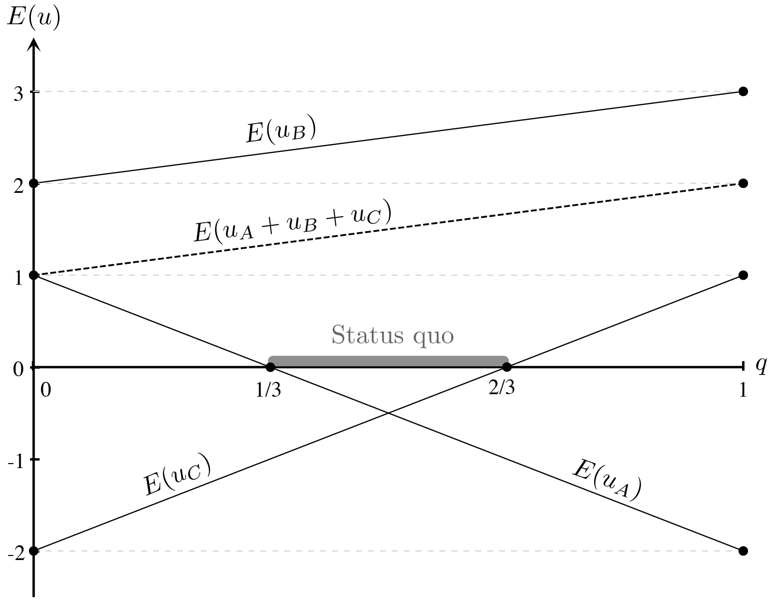

Figure 2: Expected utilities from the metro project. When q is in the grey interval, expected utility of A and C is below zero so the majority votes for status quo.

status quo. For voter i, the lower-case letters (m_i and t_i) represent the associated utility from M and T, while the utility from S is normalized to zero. The initial collective choice is either *Yes* to start construction or *No* to shelve the project. After *Yes*, an antiquity is discovered with probability q. The discovery triggers a second vote between *Continue* to finish metro and *Stop* to preserve antiquity.

Action *Yes* leads either to M or, with probability q, a choice between M and T. If option T arises, it can be ignored, in which case the project is completed as planned, or it can be exercised at the discretion of a majority, in which case the project is terminated. Inaction *No* leads to S.

Suppose that the type distribution (over tuples $(t_i, m_i) \in \mathbb{R}^2$) is such that for at least half the population,

$$t_i > m_i$$

and, for at least half the population,

$$m_i > 0$$

(a majority prefers the metro over the status quo, and another majority prefers the antiquity over the metro). Also

$$\sum_i t_i > \sum_i m_i > 0$$

(across the population, the metro yields a positive total utility, and the antiquity yields a greater total utility than the metro).[10] Therefore, from a social standpoint, the metro is desirable, and the antiquity is a valuable option if found. By assuming that, in pairwise comparisons, a majority sides with the social planner, we bias the game against the

[10]Summation is only used for ease of notation; the type support could be a continuum.

possibility of an inefficient decision on the project.

It is crucial that the majority that favors M over S can have a different composition from the majority that prefers T to S. If a majority of citizens individually preferred both M and T to S, the initial vote would clearly be in favor of the project. Since we are interested in the possibility that a majority opposes building the metro, we rule this out: less than half the population prefers M and also T to S. In order to find the expected payoff from the project for an agent, recall that the antiquity by assumption, if found, has majority support over the metro, hence the project stops. This means that the project is a lottery where i receives utility of m_i with probability $1-q$ and t_i with probability q. The utility from the status quo is assumed to be zero, so i supports the project if

$$E(u_i) = (1-q)m_i + qt_i \geq 0,$$

and is against the project otherwise.

To achieve efficiency, a social planner who wants to maximize total utility has a simple problem to solve. If an antiquity is found, he should choose the option with the highest total utility. Since $\sum_i t_i > \sum_i m_i$, the maximum expected utility from the project is given by

$$(1-q)\sum_i m_i + q \max\left\{\sum_i m_i \ , \sum_i t_i\right\} = (1-q)\sum_i m_i + q\sum_i t_i$$

where the expression $\max\{\sum_i m_i, \sum_i t_i\}$ reflects that the antiquity is an option that does not need to be exercised. Then, the social planner's implied decision rule for starting the project

$$(1-q)\sum_i m_i + q\sum_i t_i \geq 0$$

is similar to that of an individual agent, with the difference that the social planner uses total, instead of individual, expected utility to make the decision.

All that is required to stop the project is that, for a majority of voters, its expected utility is less than zero. Our simple example with three type suggests that this scenario does not require a very particular type distribution. In Appendix A, we give a more detailed discussion and derive sufficient conditions both for a majority choice pro and contra construction. These are, respectively, a tendency to regard the worst outcome as too bad to be compensated by the best outcome (a property we call upward-convergent payoffs) and the opposite tendency (downward-convergent payoffs).[11] In the first case, the fact that starting the project and maybe finding an antiquity opens the door to a new

[11] In the three-types case (equal mass), one can easily construct further numerical examples where the metro is rejected. Leaving the ranking of outcomes the same for A, B, and C, the necessary and sufficient condition on payoffs is that $m_A t_C < t_A m_C$ (this follows as a special case from Proposition A3 in Appendix A, which covers general type distributions). That is, among the pivotal types, the product of negative payoffs (antiquity for A, metro for C) must dominate the product of positive payoffs (metro for A, antiquity for C). Under these conditions, the project loses when $m_A/(m_A - t_A) < q < m_C/(m_C - t_C)$, which is a nonempty interval (given $m_A t_C < t_A m_C$). But one can always make the project efficient by raising B's payoffs enough.

alternative is a bad thing from the point of view of many citizens.[12] In the second case, more possibilities are a good thing. Moreover, with symmetric distributions, such as a uniform distribution on a circle, the efficient choice will never be rejected, because the majority must agree with the average type (social planner).

Can commitment solve the problem?

What is striking about the metro game is that a socially superior option (starting and finishing the metro) is defeated in the presence of uncertainty, even though it is always in the collective power of the voters to implement it. The problem is, of course, that voters expect a majority to force abandonment of the project in case an antiquity is discovered. Since reserving the option to do so leads to an undesirable outcome, it is natural to expect that a commitment to finishing the metro line would allow the project to go forward.

We refer to a complete plan of action, at all points where the population decides, as a policy. In the three-types example, A and B should both favor the policy "build the metro, no matter what" to the status quo, irrespective of q. However, this is not a satisfactory way to think about voting with commitment. The alternative policy "start the metro, but stop if an antiquity is found" would have majority support from B and C over the former. Unless some policy can beat *all* possible policies, it is not clear how voters could agree which policy to pit against the status quo. We run into the classic Condorcet cycle problem.

It turns out that the Roman metro game will *always* exhibit a cycle in voting over policies *if* the project is defeated in step-by-step voting without commitment. The converse is also true, so that there is a formal equivalence between failure of the project and a cycle when commitment is possible. Furthermore, we will show in the paper that the relationship is fundamental and holds for a general class of voting games. In the Roman metro game, commitment expands the initial space of alternatives from two (Y and N) to three: (YM) "start and finish the metro, whether or not an antiquity is found," (YT) "start, but preserve the antiquity if found" and (N) "do nothing."[13]

Proposition 1. *In the Roman metro game, a majority opposes the project if and only if there exists a Condorcet cycle over the set of all policies.*

Proof. If a Condorcet cycle exists, then it has to take the form that (YM) is majority-preferred to (N), (YT) is majority-preferred to (YM) and (N) is majority-preferred to (YT). The reason is that (YM) delivers metro with certainty, and by assumption a majority prefers metro to doing nothing. Moreover, (YT) is a lottery between metro and antiquity, which is majority-preferred to metro because a majority prefers antiquity to metro. The only way to get a Condorcet cycle is then for (N) to beat (YT) in majority voting. But this is precisely the choice voters make in the Roman metro game when they

[12] This is almost, but not quite, the same thing as risk aversion, which would be to refuse a fair gamble with an equal upside and downside (however, the downside weighs greater, once preferences are taken into account). Upward-convergent payoffs are such that the downside is large relative to the upside, so that the gamble is not a fair one.

[13] We do not distinguish between the policies "do nothing, but if project starts finish the metro" and "do nothing, but if project starts preserve the antiquity".

forego the project because $(1-q)\,m_i + qt_i \leq 0$ for a majority. Conversely, if a majority opposes the project, then (N) is majority-preferred to (YT). Because (YT) is majority-preferred to (YM), and (YM) is majority-preferred to (N) by the assumptions of the game, we have a Condorcet cycle. □

In our three-types example, we found that the majority votes against the project if $q \in (1/3, 2/3)$. With the probability of finding an antiquity in this interval, A and C oppose the metro project. When is there a Condorcet cycle between the policies (YM), (YT) and (N)? When (N) beats (YT), that is, $(1-q)\,m_i + qt_i \leq 0$ for a majority. This majority must consist of A and C, since B prefers both metro and antiquity to the status quo. But A and C both satisfy the inequality when $q \in (1/3, 2/3)$. Thus, the requirement for the Condorcet cycle to exist is the condition for a majority to oppose the project.

It turns out then that it is not simply the inability to commit that causes the project to be defeated. If commitment were possible, and there existed a Condorcet winner among the feasible policies, a majority would support construction of the metro even without commitment, that is, in voting on immediate actions. Whenever a majority opposes construction, there is no Condorcet winner among the feasible policies, so one cannot guarantee a better outcome by committing.

The simple example presented above suggests a relationship between undesirable equilibria in voting without commitment, and Condorcet cycles with commitment. As we show next, this relationship is rather general. Namely, in a large class of voting games with uncertainty, of which the Roman metro game is a special case, commitment to future actions cannot help voters avoid undesirable outcomes that result from voting without commitment.

3 General voting game

Consider a voting game with T periods and an odd number N of voters. Each period starts with a publicly known state θ_t belonging to some space Θ_t, which contains all the necessary information about past decisions and observations.

At each period t, a collective decision must be made from some binary set $A(\theta_t) = \{\underline{a}(\theta_t), \bar{a}(\theta_t)\}$. This choice, along with past choices and states, determines the distribution of the state at the next period. Formally, each Θ_t is associated with a sigma algebra Σ_t to form a measurable space, and θ_{t+1} has a distribution $F_{t+1}(\cdot|a_t, \theta_t) \in \Delta(\Theta_{t+1})$ given by a conditional probability system $F_{t+1}(\cdot|\cdot)$.[14]

For example, the state θ_t may represent the probability distribution of some unknown but payoff-relevant parameter $\hat{\theta}$, given the information accumulated until period t. The state θ_{t+1} then includes any new information accrued between periods t and $t+1$ about the value of $\hat{\theta}$, and such information depends on the action taken in period t.[15] In the Roman metro example, if the city decides to start the construction of a metro line in period 1,

[14] See for example Durrett [4] for a formal definition of these objects.

[15] In this case, $\Theta_t = \Delta(\Theta)$ for all t's, where Θ is the (finite, say) parameter space containing $\hat{\theta}$ and $\Delta(\Theta)$ is the set of distributions over that set. The sequence $\{\Sigma_t\}$ of sigma-algebras forms a filtration, i.e., is such that $\Sigma_{t'}$ is finer than Σ_t for all $t' \geq t$.

some ruins may be discovered with positive probability, which affects θ_2. If the city does not undertake construction, nothing is learned and θ_2 contains no further information about the existence of ruins. The state θ_t can also include a physical component, such as the current stage of a construction.

Let $\Theta = \cup_{t=1}^{T} \Theta_t$ denote the set of all possible states and $A = \cup_{\theta \in \Theta} A(\theta)$ denote the set of of all possible actions. Each voter i has a terminal payoff $u_i(\theta_T)$, which depends on all past actions and shocks, as captured by the terminal state θ_T. A *policy* $C : \Theta \to A$ maps at each period t each state θ_t into an action in $A(\theta_t)$. Similarly, a *voting strategy* for voter i is defined by a policy C^i, which describes his voting decisions.

Given a policy C and a state θ_t, voter i's expected payoff, seen from period t, is

$$V_t^i(C|\theta_t) = E[u_i(\theta_T)|\theta_t, C].$$

We assume for simplicity that for each period t, state θ_t, and policy C, each voter has a strict preference for one of the two actions in $A(\theta_t)$. That is, we rule out situations in which

$$V_t^i(\underline{a}(\theta_t) \circ C|\theta_t) = V_t^i(\bar{a}(\theta_t) \circ C|\theta_t)$$

for some i, where $\underline{a}(\theta_t) \circ C$ denotes the policy equal to C on $\Theta \setminus \{\theta_t\}$ and equal to $\underline{a}(\theta_t)$ for θ_t, with a similar definition for $\bar{a}(\theta_t) \circ C$.

Definition 1 (Voting Equilibrium). *A profile $\{C^i\}_{i=1}^N$ of voting strategies forms a* Voting Equilibrium in Weakly Undominated Strategies *if the following conditions hold*

- *The resulting collective decision Z satisfies $Z(\theta_t) = a \in A(\theta_t)$ if and only if the cardinality condition $|C^i(\theta_t) = a| \geq N/2$ holds.*
- $C^i(\theta_t) = \arg\max_{a \in A(\theta_t)} V_t^i(a \circ Z|\theta_t)$

The first condition describes simple majority voting: at each time, society picks the action that garners the most votes. The second condition corresponds to the elimination of weakly dominated strategies. At each period t, voter i, taking as given the continuation of the collective decision process from period $t+1$ onwards that will result from state θ_{t+1}, votes for the action that maximizes his expected payoff as if he were pivotal. Because indifference is ruled out and the horizon is finite, this defines a unique voting equilibrium, by backward induction.

Proposition 2. *There exists a unique voting equilibrium.*

4 Commitment and voting cycles

Let Z denote the equilibrium policy. Given two policies Y and Y', say that Y dominates Y', written $Y \succ Y'$, if there is a majority of voters for whom $V_1^i(Y|\theta_1) > V_1^i(Y'|\theta_1)$. A voting cycle is a finite list of policies $Y_0, \ldots, Y_K$ such that $Y_k \prec Y_{k+1}$ for all $k < K$, and $Y_0 = Y_K$. A policy X is a Condorcet winner if for any policy Y, either $X \succ Y$, or X and Y induce the same distribution over Θ_T.

Theorem 1. *The following statements hold:*

i) If there exists Y such that $Y \succ Z$, then there is a Condorcet cycle that includes Y and Z.

ii) If there exists a policy X that is a Condorcet winner among all policies, then X and Z induce the same distribution over Θ_T.

An immediate consequence of the theorem is the following equivalence: a Condorcet winner exists if and only if the policy Z is undominated.

The proof, below, works as follows: if Y is different from Z, then Y takes an action somewhere that the majority opposes. This allows us to construct a sequence of policies, starting from Y, where we switch actions to those preferred by a majority, thus always defeating the previous policy, until we recover Z, which was defeated by the original alternative Y. The proof collects states according to the (finitely many) winning coalitions, because individual states may (and often do) have zero probability. Because we group states by all possible majorities, and then switch actions for each majority, the switch is guaranteed to have the support of that particular majority. Proceeding by backward induction on the decision tree, this sequence of transformations recovers the equilibrium policy Z. Thus, we get a Condorcet cycle if, initially, $Y \succ Z$.

By way of illustration, consider the Roman metro game with parameter values, which were characterized earlier, such that i) the equilibrium outcome, Z, is the status quo, and ii) a majority prefers commitment to building the metro line, Y, over the status quo. In order to uncover the cycle predicted by Theorem 1, consider the majority (consisting of types B and C, in the notation of Section 2) which favors preservation if an antiquity is discovered. That majority prefers, over commitment to building the metro line, the plan of action Y_1 in which, conditional on digging and finding antiquity, the antiquity is preserved. Thus $Y_1 \succ Y$. However, Y_1, in turn, is defeated by the status quo Z because types A and C prefer not digging at all rather than having to preserve the antiquity in the case of a discovery. We thus obtain the cycle $Z \prec Y \prec Y_1 \prec Z$.

Proof. Consider any policy Y. Let $\mathcal{S}$ denote the set of coalitions with at least $N/2$ voters. For each θ_t, $a \in A(\theta_t)$ and policy X, let $S(a|\theta_t, X)$ denote the set of voters who strictly prefer a to the other action in $A(\theta_t)$, given the current state θ_t and given that the continuation policy from $t+1$ onwards is X. The set Θ_T can be partitioned into $A_T \cup (\cup_{S \in \mathcal{S}} B_T(S))$, where $B_T(S) = \{\theta_T : Z_T(\theta_T) \neq Y_T(\theta_T) \text{ and } S(Z_T(\theta_T)|\theta_T, Z)) = S\}$ and A_T consists of all remaining sets. In words, $B_T(S)$ consists of all the states in period T for which the set of voters who strictly prefer the action prescribed by Z over the one prescribed by Y is equal to S.[16] A_T consists of all the states for which Y_T and Z_T coincide.[17] We will index the coalitions in $\mathcal{S}$ from S_1 to $S_{\bar{p}}$, where $\bar{p}$ is the cardinal of $\mathcal{S}$. Consider the sequence of policies $\{Y_T^p\}_{p=1}^{\bar{p}}$, defined iteratively as follows:

- Y_T^1 is equal to Y for all states except on $B_T(S_1)$, where it is equal to Z.

[16]In particular, $Y_T(\theta_T) \neq Z_T(\theta_T)$ for all those states.

[17]By definition of Z, the set of voters who prefer Y over Z forms a minority, so A_T and $\cup_{S \in \mathcal{S}} B_T(S)$ exhaust all states in Θ_T.

- For each $p \in \{2, \ldots, \bar{p}\}$, Y_T^p is equal to Y_T^{p-1} for all states except on $B_T(S_p)$, where it is equal to Z.

By construction, $Y_T^1 \succeq Y$ because the policies are the same except on a set of states where a majority of voters prefer Z (and, hence, Y_T^1) to Y. Moreover, the preference is strict if and only if $B_T(S_1)$ is reached with positive probability under policy Y: $Y_T^1 \succ Y \Leftrightarrow Pr(B_T(S_1)|Y) > 0$. If $Pr(B_T(S_1)|Y) = 0$, $Y_T^1 = Y$ with probability 1.

Therefore, either Y and Y_T^1 coincide, or $Y_T^1 \succ Y$. Similarly, $Y_T^p \succeq Y_T^{p-1}$ for all $p \leq \bar{p}$, and $Y_T^p \succ Y_T^{p-1}$ if and only if $Y_T^p \neq Y_T^{p-1}$ with positive probability. This shows that

$$Y_T^{\bar{p}} \succeq \cdots \succeq Y_T^1 \succeq Y,$$

and at least one inequality is strict if and only if the set of states in Θ_T over which Z_T and Y_T are different is reached with positive probability under Y. By construction, $Y_T^{\bar{p}}$ coincides with Z on Θ_T: $Y_T^{\bar{p}}(\theta_T) = Z(\theta_T)$ for all $\theta_T \in \Theta_T$.

We now extend the construction by backward induction to all periods from $t = T-1$ to $t = 1$. For period t, partition Θ_t into $A_t \cup (\cup_{S \in \mathcal{S}} B_t(S))$, where A_t consists of all θ_t's over which Y_t and Z_t coincide, and $B_t(S) = \{\theta_t : Z_t(\theta_t) \neq Y_t(\theta_t) \text{ and } S(Z_t(\theta_t)|\theta_t, Z)) = S\}$. That is, $B_t(S)$ consists of all states in Θ_t for which the set of voters who strictly prefer the action prescribed by Z over the one prescribed by Y, *given that Z is used for all subsequent periods*, is equal to S.[18] Y_t^p is defined inductively as follows, increasing p within each period t, and then decreasing t: for each t,

- For $p = 1$, Y_t^1 is equal to $Y_{t+1}^{\bar{p}}$ for all states, except on $B_t(S_1)$, where it is equal to Z.

- For $p > 1$ Y_t^p is equal to Y_t^{p-1} for all states, except on $B_t(S_p)$ where it is equal to Z.

By construction, $Y_t^{p+1} \succeq Y_t^p$ for all t and $p < \bar{p}$ and $Y_t^1 \succeq Y_{t+1}^{\bar{p}}$ for all t. Moreover, the inequality is strict if and only if the policies being compared are not equal with probability 1 on the set of states reached by either of them.

Finally, observe that $Y_1^{\bar{p}} = Z$. Let $\{Y_k\}_{k=1}^K$, $K \geq 1$, denote the sequence of *distinct* policies obtained, starting from Y, by the previous construction, iterating from $t = T$ and $p = 1$ down to $t = 1$ and $p = \bar{p}$.[19]

If $Y \neq Z$ with positive probability, then $K \geq 2$. Moreover,

$$Y = Y_1 \prec Y_2 \cdots \prec Y_K = Z.$$

Therefore, we get a voting cycle if $Z \prec Y$, which concludes the proof of part i).

Since Z can never be defeated without creating a cycle, we can characterize a Condorcet winner out of all policies, if it (they) exists, and ii) follows. $\square$

[18] Again, by definition of Z, there cannot be a majority who prefer Y_t over Z_t, given the continuation policy $\{Z_t'\}_{t' \geq t}$, so $A_t \cup (\cup_{S \in \mathcal{S}} B_t(S)) = \Theta_t$.

[19] We call two policies *distinct* if they induce different distributions over Θ_T. Policies that differ only at states that are never reached are not distinct.

5 Applications

While the Roman metro game illustrated how a majority may favor clearly inefficient choices that commitment to future actions cannot solve, Theorem 1 says nothing directly about efficiency. It is mute on why a particular policy is majority-preferred to the equilibrium, if such is the case. In the metro game, majority preference follows total utility by assumption, creating a link between utilitarian inefficiency in equilibrium and a Condorcet cycle over policies. However, since there is, in general, no relationship between rankings by total utility and majority preference, there is none between efficiency in this sense and cycles either.

Matters are different if efficiency fails in the stronger Pareto sense. In our context, we say that a policy is strictly Pareto-dominated if there exists another policy that all voters strictly prefer to the initial policy. If a Pareto improvement over the equilibrium exists, then commitment to it surely defeats the equilibrium policy and sets in motion the machinery of Theorem 1.

Corollary 1. *If the equilibrium policy is strictly Pareto-dominated, then there is a Condorcet cycle over the set of all policies.*

Proof. Since the equilibrium is strictly Pareto-dominated only if there exists a policy Y that yields a higher expected payoff for all voters, hence is majority-preferred to the equilibrium policy, Theorem 1 applies immediately. □

Instances where Pareto inefficient decisions are caused by the logic studied in this paper arise naturally in many situations, as the next example illustrates.

Example 1. The following is a game of social experimentation where collective decisions may lead to Pareto inefficiency: for all participants, another course of action could guarantee higher payoffs. In a familiar job market situation, an economics department considers flying out a candidate for a macro position, who presents herself as both a macro and a labor economist.[20] It is not clear yet where her primary interest lies, but if the flyout is scheduled, it will reveal the candidate's true field. In effect the faculty decides whether to "experiment" with a new candidate (N). In case the department cannot agree on the flyout, the status quo candidate (S) will get the position; however, there is a shared feeling among the faculty that more information should be gathered before making the offer. Hence, there is an inherent preference for a flyout.

Once N's field is known, the faculty will vote on whether to hire her or S. The politics of this are as follows. To a third of the faculty (group A), it is important to hire a candidate who will exclusively work on macroeconomics; these members would prefer to make an offer to S over a labor economist. Another third (group B) is already convinced about N and willing to make an offer to N over S, regardless of whether she is macro or labor economist. The remaining third (group C) will vote for N only if she is a labor economist, and otherwise prefer S. Figure 3 displays this game. The utility from the status quo

[20]This example was suggested by Jeff Ely.

candidate is normalized to zero, and the flyout is inherently desirable in that hiring S yields a greater payoff after N's flyout. The total utility from inviting the new candidate and making an offer is always positive; furthermore, there is always a majority that would support an offer to N.

Figure 3 goes here.

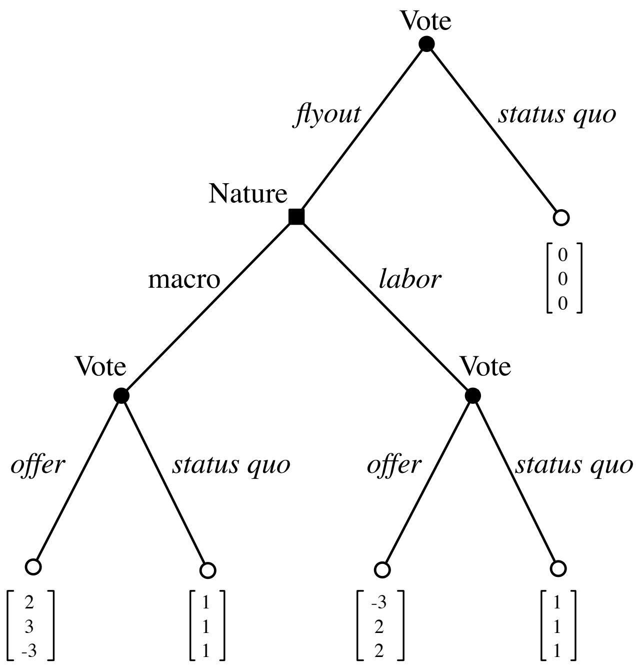

Figure 3: Job Market Game

Still, groups A and C both refuse to fly her out, for different reasons. A fears that N is found to be a labor economist, and B and C then align to make her an offer. C worries that N could turn out to be a macroeconomist and would then be supported by A and B. Hence, the candidate may fail to get the flyout, even though, regardless of what the flyout reveals, a majority would vote to make her an offer, and the benevolent department head would do the same. Moreover, in this case the choice not to fly the candidate out is Pareto inefficient, since a strictly superior outcome can be guaranteed for everyone (hiring S is always possible and valued more after the flyout). In these circumstances, commitment is a tempting solution. For instance, commitment to only hire a labor economist has majority support over the status quo, because it results in a lottery between the status quo candidate and a labor economist, whom a majority prefers. This means, however, that Theorem 1 applies: since there is a policy that defeats the status quo, voting on policies must be subject to a Condorcet cycle.[21]

[21]The cycle works as follows. Take, for example, the Pareto improvement over the status quo: flying N out, but hiring S in any case. This policy defeats the status quo, but is itself beaten by a policy to make an offer to N only if she is a labor economist, which is preferred by B and C. This policy, however, is rejected by a majority in favor of hiring N in any case, since this yields better outcomes for A and B. But flying N out and hiring her cannot defeat the status quo, since the expected payoff for A and C is negative.

An interesting aspect of the job market game is that, while decision makers have no clear way out of an inefficient choice, the candidate could fix the problem by revealing her type. By positioning herself clearly as a macroeconomist or a labor economist, she can remove any uncertainty, in which case the department will invite her. This rationalizes the advice job candidates often get from their advisors to avoid mixing fields. The corresponding action in the Roman metro game would be to locate antiquities prior to deciding on the project, so that the uncertainty is resolved. Unfortunately, this is not feasible in practice.

Example 2. In this example we discuss the relation between our framework and the social experimentation game studied by Fernandez and Rodrik [5]. Fernandez and Rodrik study a two-stage game where, at the first stage, the population decides whether to institute a trade reform or not. After the decision is made, winners and losers from the reform are realized prior to second-stage voting on whether to continue the reform.[22] A much simplified version of the original game is depicted in Figure 4. Given a status quo, a reform can be implemented by majority vote in either or both of two periods. Reform imposes a (sunk) cost c on each individual that must be borne once, regardless of the duration of the reform. There are three types of voters (A, B, and C). One of the types is randomly (with uniform probability) chosen as the sole winner from the reform. The winner gets a payoff of g per period for the duration of the reform, while all others lose l per period. If the reform is implemented in the first period and continued in the second, we call it a long-term reform, whereas if it is revoked in the second period, it is a short-term reform.

Figure 4 goes here.

Fernandez and Rodrik argue that the majority has a tendency to prefer the status quo in the following sense: if commitment to status quo in both periods is preferred to a commitment to long-term reform, then status quo also beats starting the reform in every period without commitment because, when the status quo is kept, no new information arrives, and there is no reason to reconsider the initial decision. Thus, a vote for the status quo is never followed by a vote for reform (else, reform should have been implemented immediately). In contrast, even when a commitment to long-term reform defeats a commitment to the status quo (true if gains for the winner are large enough), it is possible that the reform is not begun without commitment, since voters anticipate that it will be short-lived, given that two out of three types find out they are losers from the reform and have an incentive to end it in the second period.

The latter case is of interest because Theorem 1 states that there exists a Condorcet cycle over policies under these conditions. The status quo survives in equilibrium without commitment if $g < 2l + 2c$, so that the expected payoff from short-term reform (the

[22] Voters also participate in the labor market where they choose sectors. Our presentation here only captures the voting behavior, taking the labor market equilibrium as given. An issue raised by Ciccone [3] regarding their conclusions is related to voters' actions outside of the voting booth (that is, in the job market) and does not concern us here. Also, we assume that the winner is drawn after the reform is implemented, thus a different winner may arise after first-period and second-period reforms. This is done purely for simplicity.

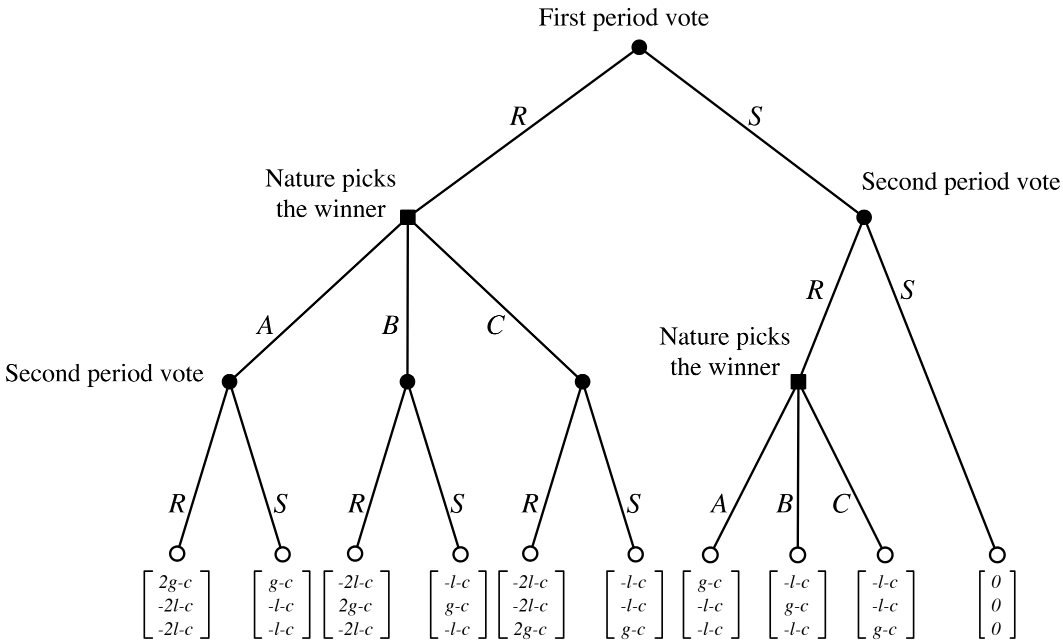

Figure 4: Reform game from Fernandez and Rodrik [5] for three voters and one winner from the reform.

reform is ended the moment two losers are identified) is negative for every type. The expected payoff from long-term reform is simultaneously positive for each type, provided that $g > 2l + 1.5c$. Other possible commitments, to short-term reform or to delayed reform in the second period, yield identical negative expected payoffs. Among these, long-term reform is a Condorcet winner, which seems to contradict Theorem 1.

How can we reconcile the above with Theorem 1? In Theorem 1, the Condorcet cycle over policies is constructed on the full set of feasible policies: any plan of action that conditions on states where voting takes place is under consideration. While natural in many contexts, such as the Roman metro game, this may be too demanding in others. In Fernandez and Rodrik, it makes little sense to let policies depend on the identity of the winner and losers. It is the anonymity that allows the commitment to long-term reform to be a Condorcet winner. But if we relax the anonymity assumption, and consider all policies that can be formulated, then long-term reform is indeed no longer a Condorcet winner. It is defeated by the following policy: start with reform initially, and revoke it only if A is the winner. Note that B and C strictly prefer this policy to unconditional long-term reform, since they will be losers in that particular state of the world. Commitment to long-term reform except when A is the winner is beaten by commitment to long-term reform unless A or B is the winner. This differs from the previous policy only when B is the winner, and in that state of the world A and C benefit from ending the reform, so their expected payoffs increase. Now that A and B are at best temporary winners, both prefer the status quo (recall that the expected payoff from short-term reform is negative). So we have a cycle: unconditional commitment to long-term reform defeats the status quo and unconditional commitment to temporary reform, but loses to long-term reform unless A is the winner, which in turn loses to long-term reform unless A or B is the winner, which in turn loses to commitment to the status quo.

Example 3. A natural application of our results concerns the stability of political regimes which are disliked by a majority of their constituents. In particular, the logic that we identified captures a possible explanation for the stability of dictatorships. To accommodate this application, we need a slight modification of the formal setting introduced in Section 4, which illustrates a potential extension of Theorem 1. The games covered in our formal analysis involve only binary choices. This restriction is, however, unrelated to the logic of the theorem; it was made only to ensure that majority voting has a winner at each stage. In many cases, it can be relaxed without affecting the validity of the theorem, as we illustrate with the following game.

Dictators are sometimes able to stay in power despite widespread dissatisfaction, partly because the opposition, while united in wishing to depose the current government, is fragmented by conflicting goals and mutual distrust.[23] Suppose that there are three groups: the elite, who benefits from the current dictatorship (the "status quo"), as well as secular and religious factions. A revolt against the status quo is successful if it is supported by a majority of the population.[24] Three outcomes are then possible: religious, secular, and bipartisan regimes. In the bipartisan regime, power is shared between the religious and secular groups and the elite loses its privileges. While the religious and secular groups both want to remove the dictator, they favor radically opposite regimes, following an uprising. Let us suppose that each constitutes 49% of the population so that the decisive factor is the elite, which represents the remaining 2%.[25] The elite is interested in preserving some degree of power and might side with either the religious or the secular group, depending on how the uprising unfolds. Finally, we assume that the elite supports each group with equal probability, but always prefers a religious or a secular regime to a bipartisan one, in which it loses all influence.

Figure 3 depicts the game, where payoffs are in the following order: religious, secular, and elite. While, as we described, the population has a choice between three distinct outcomes (regimes) after the uprising, we can easily cast the problem as a series of binary choices, so as to fit the framework of Theorem 1.[26]

Figure 5 goes here.

[23] The following is a variation on an idea that was originally suggested to us by Stephen Schmidt, in reference to the Arab Spring civil uprisings.

[24] Voting is here a metaphor for collective action. As we already pointed out in the context of the Roman game, our analysis is relevant whenever decisions depend on support from a critical mass, which may or may not take the form of explicit voting.

[25] This percentage is meant to capture the qualitative feature that no group can single handedly impose its most preferred regime, and that the elite can play a pivotal role. The "vote" that follows an uprising is a stylization of the struggle for power that ensues such uprising, and the percentages represent the relative strength of each group in that fight.

[26] More generally, it is possible to accommodate voting over three or more alternatives (at each stage), as long as expected payoffs are such that, for any continuation play, there exists a Condorcet winner among them. To avoid confusion with the issue of Condorcet winners across commitment policies, we have not included this possibility in the statement of Theorem 1.

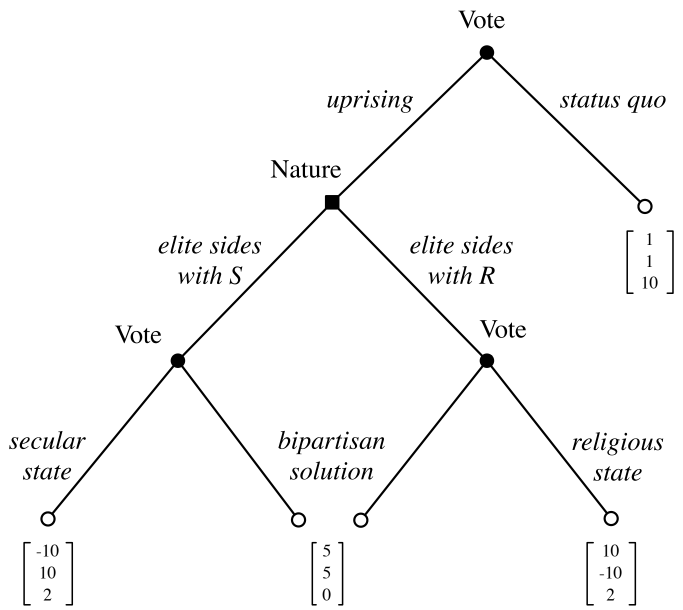

Figure 5: Dictator game

Since the elite lends its support to either the religious or the secular group, revolting amounts to a lottery between a religious or a secular regime, and has an expected payoff of zero for both of these groups - less than the status quo payoff. The dictator, whom 98% of the population would like to depose, can therefore hold on to power because of the elite's pivotal role. Yet, commitment to an uprising and subsequent bipartisan government has a higher payoff and is therefore majority-preferred to the status quo.

Theorem 1 implies that there does not exist a Condorcet winner among (commitment) policies. What is the predicted cycle, which upsets the bipartisan solution, which is preferred by 98% over dictatorship? We have already seen why the dictatorship is preferred to the fair lottery, depending on whom the elite supports, which results in a partisan (secular or religious) regime. In turn, this lottery is majority-preferred over the commitment policy which consists of imposing the religious regime if supported by the elite and the bipartisan regime if the elite supports the secular one. Such commitment is, in turn, majority-preferred (by the religious faction and the elite), to commitment to a bipartisan government. This yields the cycle predicted by our theorem. The implication is that a successful opposition to the dictator cannot necessarily be organized through an ex ante agreement. Since the uprising also fails in the absence of agreements, the unpopular dictatorship seems stable whether the population can commit or not to a specific plan of action.

6 Voting Cycles with Experimentation over an Infinite Horizon

The logic of Theorem 1 extends beyond games with a finite horizon. Indeed, while the argument used to prove the theorem relied on backward induction, a similar argument can be developed in Markovian models for which backward induction is applied to the state, rather than time. We now illustrate this idea, using the setting of Strulovici [15] and focusing on the case of three voters.[27] In that setting, time is continuous and voters, choose, at each instant, between a risky action (the "reform") and a safe action (the "status quo"). While the reform is implemented, each voter can receive good news revealing that the reform benefits him, in which case he becomes a "winner" of the reform.[28] Voters who do not receive such news, called "unsure voters", become more pessimistic about the value of the reform and eventually want to abandon it and return to the status quo, if they still form a majority at that time. Owing to the Markovian structure of the model, once the safe action is implemented, it is implemented forever: social experimentation stops at that time.

In the model, the utilitarian policy, denoted by Y, is ex ante strictly preferred by all voters to the equilibrium policy, Z, by Theorem 6 of Strulovici [15].[29] The equilibrium policy is characterized by cutoffs $p(0) > p(1)$ such that experimentation stops when the probability p of success for unsure voters reaches $p(k)$ if and only if k winners or less have been reached by that time (Theorem 1). Similarly, the utilitarian policy is determined by thresholds $q(k)$ (Theorem 2). From Theorem 3, $q(k) < p(k)$ for $k = 0, 1$: intuitively, social experimentation is more valuable to a social planner who takes into account the utility of winners, and also has a higher option value of experimentation than individual voters who have to share power. Suppose that $q(1) = 0$, which means that it is socially efficient to play the risky action forever from the moment that one winner has been observed. This condition is equivalent to $g \geq 3s$, by Theorem 2, and this parametric condition is imposed in what follows.

The thresholds $p(1)$ and $q(0)$ respectively solve the following equations, which come from indifference conditions:[30]

$$p(1) = \frac{\mu s}{\mu g + (g - s) + p(1)g - s} \tag{1}$$

and

$$q(0) = \frac{\mu s}{\mu g + (g - s) + 2(q(0)g - s)}. \tag{2}$$

Because $p(1)$ and $q(0)$ are strictly below the myopic cut-off s/g (Theorem 1), this implies that $q(0) > p(1)$.[31] All the above implies $p(0) > q(0) > p(1)$.

We now construct a cycle, based on the following modifications of the utilitarian policy. Let Y_1 denote the policy that coincides with the social optimum, Y, except that experimentation stops at the threshold $p(1)$ if the only winner observed by that time is Voter 1.

[27] Appendix B shows, similarly, that the inefficiency arising in the Roman metro game can be replicated in an infinite game of social experimentation.

[28] This model can be thought of as an infinite horizon extension of Fernandez and Rodrik [5].

[29] We refer the reader to the original paper for the results and notation used in this section. Theorem numbers also refer to the original paper

[30] See Equation (17), p. 964 with $N = 3$ and $k_N = 1$ and, respectively, Equation (7) on p. 947 with $N = 3$, $k = 0$, and $W(1, p) = \frac{1}{3}(\frac{g}{r} + 2p\frac{g}{r})$.

[31] If $p(1) \geq q(0)$, then we would have $p(1)g - s \geq q(0)g - s > 2(q(0)g - s)$, where the strict inequality comes from $q(0) < s/g$. This would imply that the RHS of (1) is strictly smaller than the RHS of (2) and, hence, that $p(1) < q(0)$.

That policy is strictly preferred by Voters 2 and 3 over the utilitarian policy: conditional on Voter 1 being the only winner by the time $q(0)$ is reached, the policy that maximizes Voters 2 and 3's common expected utility consists precisely in stopping experimentation at the cutoff $p(1)$ if none of them has become a winner by then and, otherwise, in playing the risky action forever. To see this, notice that the utilitarian cutoff $\tilde{q}$ for Voters 2 and 3, starting from $q(0)$, and given that any winner will cause L to be played forever, is characterized by the indifference equation[32]

$$\tilde{q}g + 2\tilde{q}\lambda\left(\frac{1}{2}\left(\frac{g}{r} + \tilde{q}\frac{g}{r}\right) - \frac{s}{r}\right) = s.$$

The first term of the LHS is the average flow payoff for Voters 2 and 3 when the risky action is played. The second term of the left-hand side is their average jump in utility if one of them becomes a winner, multiplied by the flow probability of that event. The right-hand side is the flow payoff if the safe action is chosen instead (in that case, nothing is learned, so there is no jump in utility).

After rearrangement, this equation is identical to (1) and thus yields the same cutoff.[33] Therefore, $Y_1 \succ Y$. Similarly, consider the policy Y_2 that is identical to Y_1 except that experimentation stops at $p(1)$ if only Voter 2 has become a winner by that time. This policy is preferred by Voters 1 and 3 over Y_1, by the same reasoning, so $Y_2 \succ Y_1$. Also construct Y_3 by modifying Y_2 so that experimentation stops at $p(1)$ if only Voter 3 is a winner by that time. $Y_3 \succ Y_2$.

By construction, Y_3 stops experimentation if a single winner has been observed by the time $p(1)$ is reached, as does the equilibrium policy Z. The only difference between Y_3 and Z is that Z stops experimentation earlier, at the threshold $p(0) > q(0)$, if no winner has been observed, whereas Y_3 stops experimentation at $q(0)$.

However, subject to the constraint that experimentation stops at $p(1)$ if only one winner has been observed by that time, and continues forever otherwise, the socially optimal policy is actually to stop at $p(0)$, not $q(0)$. To see this, notice that, due to the constraint, the continuation value of a voter who becomes a winner is precisely the continuation value w under the equilibrium policy, and the continuation value for unsure voters when another voter becomes a winner is the equilibrium continuation value, u.

Formally, the socially optimal cut-off $\hat{q}$ for stopping experimentation when no winner has been observed, given the constraint on continuation play, solves the indifference equation

$$3\hat{q} + 3\lambda\hat{q}\left(w(1,\hat{q}) + 2u(1,\hat{q}) - 3\frac{s}{r}\right) = 3s.$$

The first term of the LHS is the aggregate expected flow payoff for the three voters if L is played, and the second term is their aggregate jump in utility if one of them becomes a winner, multiplied by the flow probability of this happening. The right-hand side is the aggregate flow payoff if the safe action is played instead. Dividing by 3, one obtains

[32]See Equation (7) on p. 947 with $N = 2$ and $k = 0$, and $W(1,p) = 1/2(g/r + pg/r)$.

[33]It is easily checked that Equation (1) has a unique positive solution (the left-hand side is increasing in p which the right-hand side is decreasing p).

exactly the characterization of the cut-off $p(0)$.[34]

This shows that $Z \succ Y_3$ and yields the voting cycle $Z \prec Y \prec Y_1 \prec Y_2 \prec Y_3 \prec Z$.

7 Conclusion

In the Roman metro problem, citizens decide whether to construct a metro line, which has a positive value for a majority, and which yields with some probability an option that a majority considers to be still better (finding a valuable antiquity). A majority may nevertheless prefer the status quo instead. The example demonstrates that majority voting under uncertainty can lead to starkly inefficient choices.

Standard solutions to inefficient outcomes of collective decision processes include ex post redistribution of the benefits, or ways to reframe the choice problem. Compensating losers is fraught with practical difficulties in a setting where preferences are private information. Instead, we allowed for voting over policies that involve commitment. Surely, under the circumstances, a majority could agree to start the metro project, if completion was assured? Not only does commitment fail in the Roman metro game, but we presented a general theorem linking undesirable equilibria to Condorcet cycles when commitment is feasible in finite, binary voting games. This class includes many games where projects or experiments unfold over time, with new possibilities arising as a result of initial actions.

Given the scope of the theorem, its conclusion that there can be no agreement on commitment seems hard to escape. In the job market game, it was within the power of the candidate to resolve uncertainty and thereby remove the underlying problem. In some situations, such as Fernandez and Rodrik's reform game, the commitment options may be naturally limited, so that the theorem does not need to hold. However, attempts to deliberately limit the policies to be voted on are not promising, since they run into a higher-order problem. The lack of a Condorcet winner on the full set of commitments will translate into similar indecisiveness regarding what subset should be considered. In interpreting the result, it is important to remember that difficulty to agree on a policy is a different issue from whether any particular policy would be an improvement over the status quo. The cycle may arise because of an abundance, rather than a lack, of attractive options, once the option to commit removes the procedural constraints of the dynamic game. Indeed, most of our examples have this flavor.

[34]See Equation (20), p. 965, with $N = 3$ and $k_N = 1$, and replacing the factor $N - k_N - 2$ by $N - k_N$).

References

[1] S. Barbera, M. Maschler, and J. Shalev. Voting for voters: A model of electoral evolution. *Games and Economic Behavior*, 37(1):40–78, 2001.

[2] B. D. Bernheim and S. N. Slavov. A solution concept for majority rule in dynamic settings. *The Review of Economic Studies*, 76(1):33–62, 2009.

[3] A. Ciccone. Resistance to reform: Status quo bias in the presence of individual-specific uncertainty: Comment. *American Economic Review*, 94(3):785–795, 2004.

[4] R. Durrett. *Probability: Theory and Examples*. Duxbury Press, Belmont, CA, 1995.

[5] R. Fernandez and D. Rodrik. Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review*, 81(5):1146–55, 1991.

[6] W. Jack and R. Lagunoff. Dynamic enfranchisement. *Journal of Public Economics*, 90(4-5):551–572, 2006.

[7] J. S. Kelly. The Ostrogorski paradox. *Social Choice and Welfare*, 6(1):71–76, 1989.

[8] M. Messner and M. K. Polborn. The option to wait in collective decisions and optimal majority rules. *Journal of Public Economics*, 96(5-6):524 – 540, 2012.

[9] H. Nurmi and D. G. Saari. Connections and implications of the Ostrogorski paradox for spatial voting models. *Collective Decision Making*, 43:31–56, 2010.

[10] E. M. Penn. A model of farsighted voting. *American Journal of Political Science*, 53(1):pp. 36–54, 2009.

[11] D. W. Rae and H. Daudt. The Ostrogorski paradox: a peculiarity of compound majority decision. *European Journal of Political Research*, 4(4):391–398, 1976.

[12] K. Roberts. Condorcet cycles? a model of intertemporal voting. *Social Choice and Welfare*, 29(3):383–404, 2007.

[13] K. A. Shepsle. A note on Zeckhauser's "Majority rule with lotteries on alternatives": The case of the paradox of voting. *Quarterly Journal of Economics*, 84(4):705–09, 1970.

[14] B. Strulovici. Voting and experimentation. Working paper, Oxford University, 2007.

[15] B. Strulovici. Learning while voting: Determinants of collective experimentation. *Econometrica*, 78(3):933–971, 05 2010.

[16] R. Zeckhauser. Majority rule with lotteries on alternatives. *Quarterly Journal of Economics*, 83(4):696–703, 1969.

Appendix A: Formal Results for the Metro Game

In this appendix, we consider properties of the type distribution which guarantee either that the metro project fails or succeeds in majority voting. By Theorem 1, these conditions are also sufficient for a Condorcet cycle on policies to exist or not exist in the Roman metro game. The properties are closely related to risk tolerance. The necessary and sufficient condition for a majority to oppose the metro project is similar.

Let u_i^1 denote type i's utility from the top-ranked alternative, u_i^2 the utility from the mid-ranked alternative, and u_i^3 the utility from the bottom-ranked alternative.

Definition: Type i's payoffs are *upward-convergent* if $u_i^1 - u_i^2 \leq u_i^2 - u_i^3$. Type i's payoffs are *downward-convergent* if $u_i^1 - u_i^2 \geq u_i^2 - u_i^3$.

That is, with upward-convergent payoffs (UCP), the top-ranked outcome improves on the mid-ranked outcome by less than the mid-ranked outcome improves on the bottom-ranked outcome. In other words, moving from the worst to the best outcome, one gains smaller increments of utility. With downward-convergent payoffs (DCP), the situation is opposite. Moving from the worst to the best outcome, one gains larger increments of utility.

Though such properties might conceivably hold for the entire population, for our purposes only a critical subset of types needs to have them. Let i be a *pivotal type* if $u_i^2 = 0$ or, expressed differently, if $\max\{m_i, t_i\} \geq 0 \geq \min\{m_i, t_i\}$. That is, i neither prefers both M and T to the status quo, nor prefers the status quo to both M and T. This means that i's preference between the lottery that yields $(1-q)\,m_i + qt_i$ and the status quo depends on the magnitude of q.[35] A pivotal type satisfies UCP if and only if $\max\{m_i, t_i\} \leq -\min\{m_i, t_i\}$, and DCP if and only if $\max\{m_i, t_i\} \geq -\min\{m_i, t_i\}$.[36] Starting at the status quo, with UCP, I lose more from switching to my least-preferred option than I gain from switching to my most-preferred option. With DCP, it's the opposite. Therefore, a pivotal type prefers 0 to an even-odds (fifty-fifty) gamble between M and T under UCP, but prefers the gamble under DCP. (This is the connection with risk tolerance.)

With UCP, the support of the type distribution is the darker shaded area in the left panel of Figure 6 (if UCP only applies to pivotal types, it is the whole shaded area).

Figure 6 goes here.

Our introductory example in the text satisfied UCP There, we had $\left(u_A^1, u_A^2, u_A^3\right) = (1, 0, -2) = \left(u_C^1, u_C^2, u_C^3\right)$ and $\left(u_B^1, u_B^2, u_B^3\right) = (3, 2, 0)$. If we were to adjust these payoffs to obtain DCP for the pivotal types A and C, namely such that $(m_A, t_A) = (2, -1)$ and $(m_C, t_C) = (-1, 2)$, it is easy to see that, at any q, either A or C (if not both) would

[35] The assumption we are maintaining throughout, that types who prefer both M and T to the status quo are in the minority, ensures that pivotal types are of consequence: at least some of them are needed to get a majority in favor of starting the metro project.

[36] One can see this directly from the definitions since, with $u_i^2 = 0$, UCP reduces to $u_i^1 \leq -u_i^3$ and DCP to $u_i^1 \geq -u_i^3$.

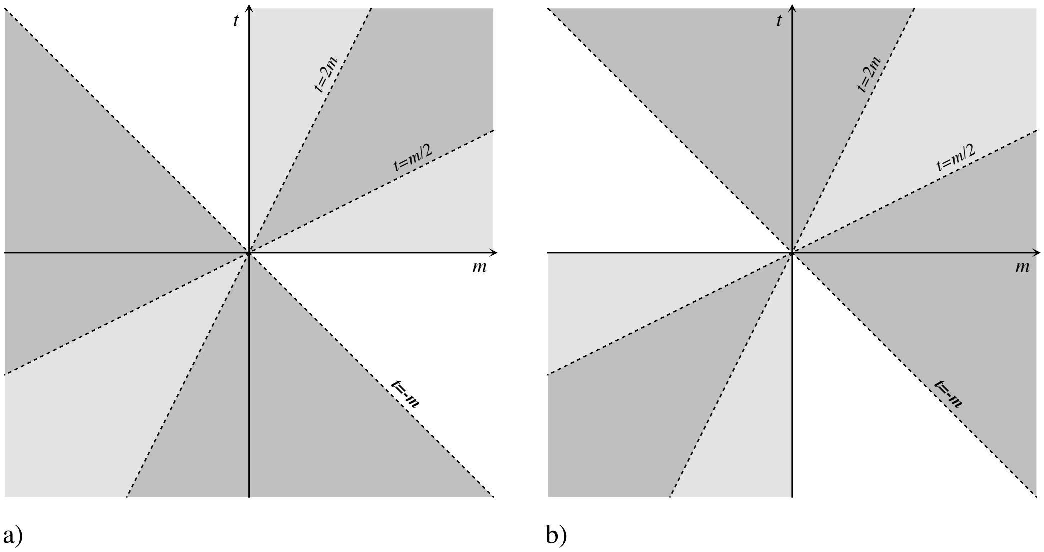

Figure 6: Type Support with UCP and DCP

support the project and form a majority with B, since

$$qm_A + (1-q)\,t_A = 2q - (1-q) \leq 0$$

only if $q \leq 1/3$, which implies

$$qm_C + (1-q)\,t_C = -q + 2\,(1-q) \geq -\frac{1}{3} + 2\left(1 - \frac{1}{3}\right) \geq 0.$$

We have two main sufficiency results. First, under UCP for pivotal types, a majority that opposes the metro project is always supported by some probability $q \in [0,1]$ that an antiquity is found. Second, under DCP, a majority always supports the project at any q.

Proposition A1. If all pivotal types have upward-convergent payoffs, then there exists a nonempty interval for q such that a majority opposes the project.

Proof. We show that the votes of pivotal types (who rank S in the middle) are sufficient to achieve a majority that opposes the project, and that there exists a q such that both groups will in fact vote against it. By assumption, types who rank S at the bottom ($m_i \geq s_i \geq 0$ and $t_i \geq m_i \geq 0$) are a minority. Those who rank S on top ($0 \geq m_i \geq t_i$ or $0 \geq t_i \geq m_i$) will certainly oppose the project, since they prefer S to anything else. Thus, getting the remaining types, who rank S in the middle, to oppose the project is enough for a majority.

Any type i opposes if $(1-q)\,m_i + qt_i \leq 0$. For someone with preference $m_i \geq 0 \geq t_i$, this is true if

$$q \geq \frac{m_i}{m_i - t_i} \equiv \underline{q}_i.$$

Similarly, someone with preference $t_i \geq 0 \geq m_i$ opposes if

$$q \leq -\frac{m_i}{t_i - m_i} \equiv \overline{q}_i.$$

With UCP for pivotal types, $\underline{q}_i \leq 1/2$ for all i such that $m_i \geq 0 \geq t_i$, since

$$\frac{m_i}{m_i - t_i} \leq -\frac{t_i}{m_i - t_i}$$

implies $q \leq 1 - q$. Similarly, $\overline{q}_i \geq 1/2$ for all i such that $t_i \geq 0 \geq m_i$. Define MST and TSM as the set of voters for whom $m_i \geq 0 \geq t_i$ and $t_i \geq 0 \geq m_i$, respectively. Then,

$$0 \leq \max_{i \in MST} \underline{q}_i \leq 1/2 \leq \min_{i \in TSM} \overline{q}_i \leq 1,$$

which means there exists a q (in an interval that contains 1/2) such that anybody who ranks S in the middle opposes the project. Given that less than half the population ranks S below both M and T, this is enough for a majority to vote against the project. $\square$

The votes align against the project because those who would rank the antiquity last (the A-s in our initial example) want to shut this alternative out, even at relatively low q that make finding the antiquity unlikely. This eliminates intervals where those who like and dislike the antiquity might both support the project, since the latter require a lower q than the former are willing to accept. In consequence, the metro line may not be built.

Figure 7 illustrates how UCP (so that the type support is the shaded area) in combination with the fact that a majority of types does not prefer both M and T to the status quo (occupies the three outlined quadrants, where at least one of m and t is negative) brings about the possibility that a majority will oppose the project. As q approaches 1/2 (the slope of the indifference line approaches $q/(1-q) = -1$), more types fall below the line, until at $q = 1/2$ a majority is sure to vote against the project.

Figure 7 goes here.

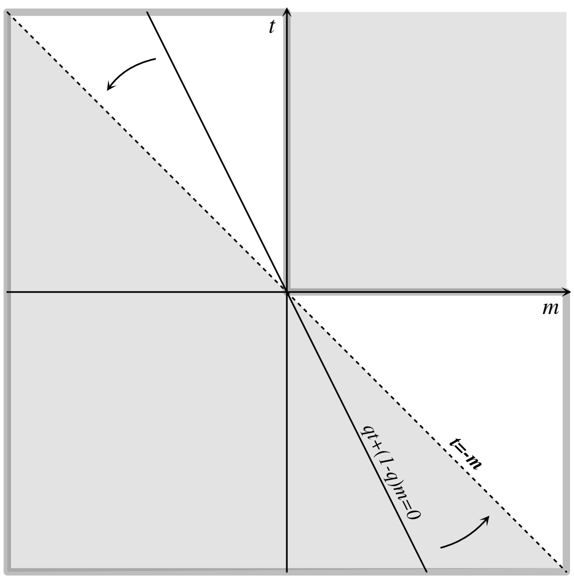

Figure 7: Project Fails with UCP at $q = 1/2$.

Proposition A2. If all pivotal types have downward-convergent payoffs, then a majority supports the project, regardless of q.

Proof. DCP for pivotal types implies

$$\frac{1}{2}m_i + \frac{1}{2}t_i \geq 0$$

for those who rank 0 in the middle, i.e. $m_i \geq 0 \geq t_i$ or $t_i \geq 0 \geq m_i$. Suppose, for the sake of contradiction, that a majority strictly opposes the project at some $q \in [0,1]$. Which types could comprise this majority? Type i would strictly oppose only if

$$(1-q)\,m_i + qt_i < 0.$$

Given DCP for pivotal types, one of (i) $0 \geq m_i$ and $0 \geq t_i$, (ii) $m_i \geq 0 \geq t_i$ and $q > 1/2$ or (iii) $t_i \geq 0 \geq m_i$ and $q < 1/2$ has to be true then. Therefore, if $q < 1/2$, types that satisfy (i) and (ii) must have a majority for there to be a majority that opposes. But then S is majority-preferred to T, which violates the assumptions. If $q > 1/2$, then types that satisfy (i) and (iii) must have a majority. Now, S is majority-preferred to M, another violation.

The only remaining possibility, $q = 1/2$, is ruled out directly by DCP for pivotal types, since all who rank S in the middle would support the project, leaving only (i) to oppose it. But (i) cannot be a majority (again, because by assumption M and T are majority-preferred to S.) Hence, there is no q that could produce a majority against the project and satisfy the assumptions as well as DCP for pivotal types. □

The votes align for the project because those who would rank the antiquity first (the C-s in our three-types example) would like to enable this alternative, even at a relatively unfavorable (low) q. This creates an interval of q-s where those who like and dislike the antiquity are both willing to support the project.

In Figure 8, we see how DCP ensures support for the project because of the assumptions that T and M are majority-preferred to the status quo. If $q > 1/2$, the indifference line is relatively flat: its slope is $-(1-q)/q > -1$. Then the shaded area within the outlined two quadrants lies completely above it. These two quadrants contain all types who prefer T to the status quo (i.e. for whom $t_i > 0$), which is a majority of the population.

Figure 8 goes here.

If $q < 1/2$, the indifference line is relatively steep with a slope smaller than -1. Then the shaded area within the two outlined quadrants lies completely above it. These two quadrants contain all types in the support who prefer M to the status quo (i.e. for whom $m > 0$), and again this is a majority.

UCP for pivotal types is not necessary for a majority to oppose the project. Suppose, in violation, $(m_A, t_A) = (2, -1)$, while $(m_B, t_B) = (5, 6)$ and $(m_C, t_C) = (-4, 1)$. There is a majority for T over M over S. The total utility from T (6) exceeds that from M (3), which exceeds that from S. Despite the fact that UCP does not hold for one pivotal type, a majority opposes the project at $q = 3/4$, since

$$(1-q)\,m_A + qt_A = \frac{1}{4}(2) + \frac{3}{4}(-1) = -\frac{1}{4} \leq 0$$

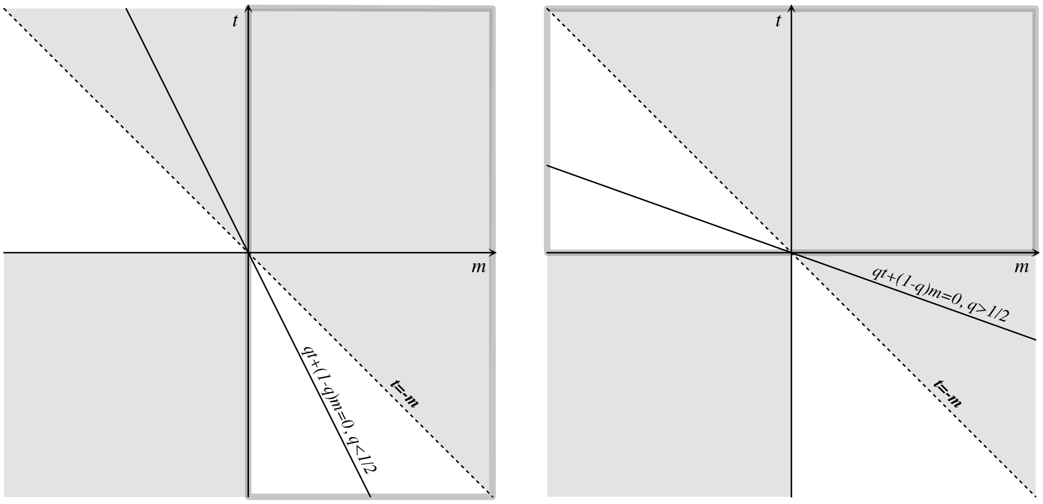

a) b)

Figure 8: Project Succeeds with DCP at all q

and

$$(1-q)\,m_C + qt_C = \frac{1}{4}(-4) + \frac{3}{4}(1) = -\frac{1}{4} \leq 0.$$

If we now increase t_C to 4, DCP is satisfied for the pivotal types. But then

$$(1-q)\,m_C + qt_C = \frac{1}{4}(-4) + \frac{3}{4}(4) = 2 \geq 0,$$

so we have a majority for the project, consistent with Proposition A2. In fact, for C to oppose the project, q would now have to be smaller than $2/5$, but no such q can induce A to oppose it, i.e. solve

$$(1-q)\,(2) + q\,(-1) = 2 - 3q \leq 0.$$

There is a majority in support at any q.

To illustrate, the left panel of Figure 9 gives a graphic representation of the three types' expected utility from the metro project in our original example.[37] Each line depicts one type's expected utility of the lottery over M and T for every $q \in [0,1]$. To defeat the project at some q, it must have negative expected utility for both pivotal types (A and C) at a common q, i.e. their lines must intersect in the negative half.

Figure 9 goes here.

UCP requires exactly that expected utility is negative at $q = 1/2$ (the negatively value of the worst outcome dominates the positive value of the best outcome). If this is true for all pivotal types, then pivotal types vote jointly against the project at least at $q = 1/2$.[38]

However, we can see that other scenarios exist where A and C both oppose the project, even though one of them gets positive expected utility at $q = 1/2$. Namely, the right panel of Figure 9 shows the payoffs for A and C in the case above, where $(m_A, t_A) = (2, -1)$ and

[37] We thank an anonymous referee for this suggestion.

[38] Similarly, DCP requires exactly that expected utility is positive at $q = 1/2$, which implies that the pivotal types' lines cannot both lie in the negative region at the same q.

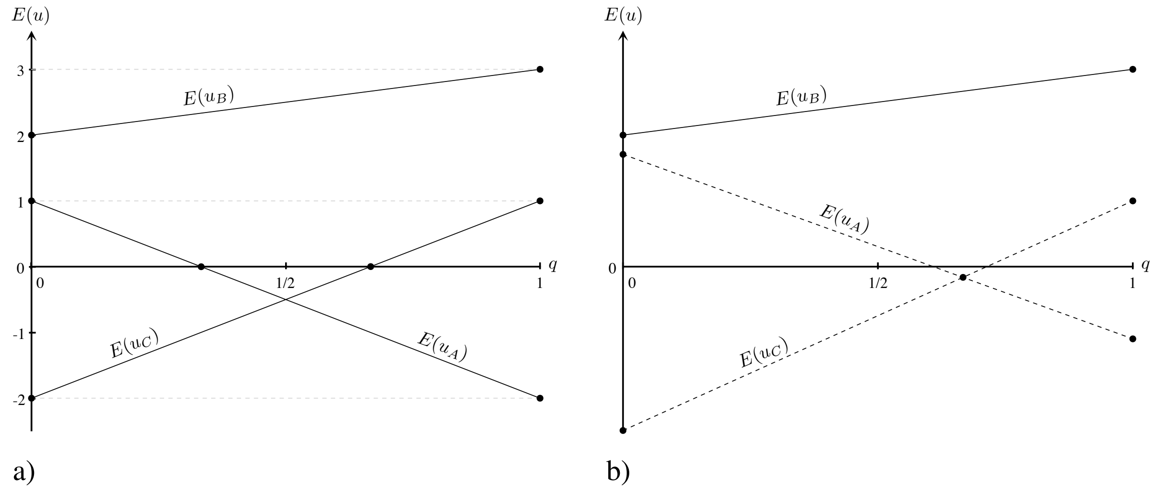

Figure 9: The left panel depicts our original example with UCP. The right panel illustrates that UCP is not necessary for the project to be defeated.

$(m_C, t_C) = (-4, 1)$. Expected utility from the project at $q = 1/2$ is positive for A; yet, there are q (larger than 1/2) at which A and C each value the project negatively. The following result generalizes the intuition from upward- and downward-convergent payoffs to provide the necessary and sufficient condition such that the metro project loses in majority voting.

Proposition A3. The metro project will be rejected for some q if and only if there exists a subset L of types, containing a majority of the population, such that (1) $\min\{m_i, t_i\} < 0$ for all $i \in L$, and (2) for any pair $i, j \in L$ such that $m_i > 0 > t_i$ and $t_j > 0 > m_j$, we have $m_i t_j < t_i m_j$. If the condition holds, then the interval of q where all members of L oppose the project is bounded below by $\max_{i \in L^m} m_i / (m_i - t_i)$ (where L^m is the nonempty set of $i \in L$ for whom $m_i > 0 > t_i$) or 0 (if L^m is empty), and is bounded above by $\min_{j \in L^t} m_j / (m_j - t_j)$ (where L^t is the nonempty set of $i \in L$ for whom $t_i > 0 > m_i$) or 1 (if L^t is empty).

Proof. Suppose there is a set L with these properties that represents a majority of the population. If neither $m_i > 0 > t_i$ nor $t_i > 0 > m_i$ for some $i \in L$, then (1) requires $\max\{m_i, t_i\} < 0$, and i votes against the project regardless of q. We show, therefore, that the remaining members of L oppose the project when q is in the specified interval, and that the interval is nonempty. If $m_i > 0 > t_i$, then i rejects the project if and only if $(1-q)\,m_i + q t_i < 0$, i.e. $q > m_i / (m_i - t_i)$. Thus, if $q > \max_{i \in L^m} m_i / (m_i - t_i)$, then all such i oppose. If $t_j > 0 > m_j$, then j rejects the project if and only if $q < -m_j / (t_j - m_j)$. Thus, if $q < \min_{j \in L^t} m_j / (m_j - t_j)$, then all such j oppose. The implied q-interval is nonempty as long as the largest $m_i / (m_i - t_i)$ from L^m is smaller than the smallest $m_j / (m_j - t_j)$ from L^t, hence if, for all i and j, with $i \in L^m$ and $j \in S^t$, $m_i / (m_i - t_i) < m_j / (m_j - t_j)$. Since $m_i > t_i$ and $m_j < t_j$, this simplifies to $m_i t_j < t_i m_j$, condition (2).[39]

Conversely, suppose no such set L can be drawn from a majority of the population.

[39] Of course, if L^m or L^t is empty, then (2) holds trivially.

Then, to construct a majority against the project, we need at least one pair of voters i and j such that $m_i > 0 > t_i$ and $t_j > 0 > m_j$, and $m_i t_j \geq t_i m_j$. In this case, however, the range of q required for i to vote against is disjoint from the range of q needed for j to vote against. Hence, only one can oppose the project, and since both are essential (otherwise, L contains a majority in the first place, contrary to assumption), the project succeeds. □

The key condition in Proposition A3, namely that $m_i t_j < t_i m_j$ for pivotal voters i and j who vote against the project, and of whom i prefers the metro and j the antiquity outcome, is similar to UCP. It says that the potential losses (t_i from antiquity for i and m_j from metro for j) outweigh the potential benefits. The sufficient condition (Proposition A1) imposed it at the individual level; the necessary and sufficient condition (Proposition A3) imposes it jointly. In our initial three-types example, where only two types are pivotal and prefer, respectively, the metro and the antiquity to all else, we can apply the result directly to derive the interval of q on which the project fails ($m_A / (m_A - t_A) = 1/3$ is the lower bound, and $m_C / (m_C - t_C) = 2/3$ is the upper bound).

Lastly, assume as in the Roman metro game that the total utility is aligned with voting majorities ($\Sigma_i t_i > \Sigma_i m_i > 0$). Then two "natural" classes of type distributions, symmetric distributions and distributions that are generated by independent draws from univariate uniform distributions, guarantee that the metro project wins the majority vote.[40]

Proposition A4. If (a) the type distribution is symmetric about a point, or if (b) valuations m and t are independently uniformly distributed (so that the type distribution is uniform on a rectangle), then a majority supports the project, regardless of q.

Proof. If a distribution is symmetric about a point, its coordinates are the population averages of m and t. According to our assumptions (that, for a majority, $m \geq 0$ and $t \geq 0$), this point has to lie in the positive quadrant.[41] The same is true of the center of a rectangular support. Clearly, a hypothetical type that occupies the center will support the project, since this type prefers both antiquity and metro to the status quo. Every line that goes through the center of a symmetric distribution or a rectangle is a median (i.e. a line that divides the space into half-spaces containing equal mass).[42] In particular, the line with slope $-(1-q)/q$ passing through the center is a median. Since this line is parallel to, and lies above, the line with slope $-(1-q)/q$ passing through the origin, which is the indifference line, a majority values the project at least as much as the center. Since the center supports the project, a majority must support it, too. □

[40] Uniform distributions on arbitrary (non-rectangular) shapes, for example triangles, need not satisfy this.

[41] Our assumptions rule out that the distribution is symmetric about zero, since there would be no strict majority preferences over outcomes.

[42] More precisely (to allow for distributions whose support consists of disjoint sets), there exists, for every median, a parallel hyperplane (median) that passes through the center and divides the population in the same way.

Appendix B: The Roman Metro Problem over an Infinite Horizon

The logic of the Roman Metro example can be extended to an infinite-horizon model. In particular, Strulovici [14] studies experimentation between several restaurants, including a Singaporean restaurant which, in addition to serving Indian cuisine, can teach "voters" about Chinese cuisine. Reinterpreting that example in the context of the Roman Metro problem, we obtain the following setting.

There are three alternatives:

- S: do not start construction
- M: build the metro, even if some ruins are discovered
- T: start construction and restore ruins if some are discovered (do not finish construction)

There are three voters, with the following preferences:

- Voter 1 cares only about the metro
- Voter 2 prefers restoring the ruins to building the metro, and building the metro to not digging at all
- Voter 3 does not care about the metro, and prefers restoring the ruins if some are found

Flow payoffs:

- The flow payoff for not starting the project is normalized to 1 for all constituents
- If the metro is built, Voters 1 and 2 get a flow payoff of 2, while Voter 3 gets 0
- If the ruins are restored, Voter 1 gets a flow payoff of -9, Voter 2 gets 2.1 and Voter 3 gets 1.1

The ex ante probability of finding some ruins is p. The analysis of Strulovici [14] implies that

Proposition B1. Absent T, M is implemented by the only Majority Voting Equilibrium (MVE). If T is added to the set of feasible alternatives and $p \in (0.1, 0.8)$, S is implemented forever by the only MVE.

Proof. In the example analyzed by Strulovici [14], Voter 3 gets a payoff $\tilde{g} = 0.1$ under M if some ruins are discovered but the metro is built nonetheless (action M). The analysis of that example is identical for all $\tilde{g} \in [0, 0.1]$, and in particular for $\tilde{g} = 0$ as is in the Roman Metro example.[43] □

[43] A possible interpretation of a positive $\tilde{g}$, in the context of the Roman Metro, is that Voter 3 gets a small positive payoff in case ruins are discovered, even the metro ends up being built. The paradox applies to that case too.