AER: Insights 2022, 4(4): 459–472
https://doi.org/10.1257/aeri.20210494

Getting Permission†

By Peicong Hu and Joel Sobel*

A manager has access to expert advisers. The manager selects at most one project and can implement it only if one expert provides support. The game in which the manager consults experts simultaneously typically has multiple equilibria, including one in which at least one expert supports the manager's favorite project. Only one outcome, the experts' most preferred equilibrium outcome, survives iterated deletion of weakly dominated strategies. We show that no sequential procedure can perform better for the manager than the experts' most preferred equilibrium and exhibit a sequential protocol that does as well. (*JEL* C72, D23, D82)

We study situations in which an individual cannot carry out a task without expert assistance. We focus on applications in which a conflict of interest may interfere with the ability of the individual to achieve his most preferred outcome but in which he can leverage competition between experts to improve his outcome.

The elements of the model are a finite set of projects, a finite set of experts, and a manager. The experts and manager have preferences defined over projects. The manager wishes to carry out a project but can do so only if an expert supports it. We are interested in the relationship between how the manager requests support and the project selected. If no expert supports any project, then the outcome is the status quo. Otherwise, the manager implements the best project consistent with the experts' approvals. Consider two alternative organizations. In the first organization, the manager asks experts to report simultaneously which of the projects they will support. In the second organization, the manager consults experts sequentially. In the first case, provided that there are at least two experts, there is always an equilibrium in which the manager receives the support needed to carry out his favorite project. If one expert supports this project, then the manager will ignore the behavior of the other experts. So it is a best reply for all of the other experts to support the manager's favorite. Sequential consultation may not work as well for the manager. In particular, if there is a project that all experts prefer to the manager's favorite, then

* Hu: Faculty of Business and Economics, University of Hong Kong (email: kerihu@hku.hk); Sobel: Department of Economics, University of California San Diego (email: jsobel@ucsd.edu). Dirk Bergemann was coeditor for this article. We thank audiences at several seminars and conferences and Snehal Banerjee, Jeremy Bertomeu, Laura Doval, Jeff Ely, Emir Kamenica, Margaret Meyer, Inés Moreno de Barreda, Ce Liu, Philip Neary, Garey Ramey, Christopher Turansick, Yuehui Wang, Joel Watson, Nicholas Wu, and referees for comments. This manuscript is a revision of a manuscript circulated under the title "Simultaneous versus Sequential Disclosure." Sobel is grateful to the National Science Foundation (award number 1757250) for research funding.
† Go to https://doi.org/10.1257/aeri.20210494 to visit the article page for additional materials and author disclosure statement(s).

sequential consultation will never provide the manager with permission to carry out his favorite project.

We want to know how the manager should organize consultation to maximize his payoff. Naïvely, the result that the simultaneous-move game includes an equilibrium that supports the manager's favorite outcome provides an answer to this question: the manager achieves his best possible outcome by consulting simultaneously. We believe that the manager-preferred equilibrium is an implausible prediction in many cases, however. We prove that experts have common preferences over equilibria (even if they have conflicting preferences over projects). The experts' preferences over equilibria are completely opposed to the preferences of the manager; that is, if the manager prefers equilibrium project x to x', then all experts prefer x' to x. (If an expert preferred project x to project x', then x' could not be an equilibrium outcome because the expert who preferred x could deviate and support x.) An equilibrium refinement (iterated deletion of weakly dominated strategies) selects the experts' preferred equilibrium.[1] Hence, the refinement rejects the manager's preferred outcome whenever another equilibrium exists. Section III presents the results. On the basis of the weak-dominance refinement (and intuition), we view the experts' preferred equilibrium as the most plausible outcome of the simultaneous-move game. This raises two questions. First, what is the value of having an additional expert? If we selected the manager-preferred equilibrium, the answer to the question is simple. Going from one expert to two experts is valuable (unless the manager's favorite task is also the initial expert's favorite task). Adding a third expert, however, has no value. When we focus on the expert-preferred equilibrium, adding experts may lead to a more attractive outcome for the manager. Our characterization implies that the manager gains by adding an additional expert if doing so makes the experts' preferred equilibrium more attractive to the manager. If x is the prediction of the simultaneous-move game with a fixed set of experts, adding an additional expert benefits the manager if there exists a project x' that both the manager and the new expert prefer to x.

The second question we study is whether an alternative organization could do better for the manager than the simultaneous game. Because our prediction for the simultaneous-move game is typically less than the manager's most preferred option, there is a chance that other procedures would work better. Furthermore, sequential consultations are the norm in some applications (e.g., obtaining a second opinion on medical procedures), so we were interested to know if there is a justification for this behavior in our setting. We identify an optimal sequential procedure (within a class of procedures that, like the simultaneous game, do not allow the decision-maker to limit the set of projects that experts can support) that does as well for the manager as the simultaneous game. No other sequential procedure can do better than it; that is, a properly designed sequential organization does at least as well as, but no better than, simultaneous consultation. We show that the following procedure performs as well as simultaneous consultation. Suppose there are K projects and the manager has strict preferences. In round r of the procedure, the manager approaches the first expert and stops if she supports one of the manager's r favorite projects. If she does

[1] The result requires an assumption that holds for generic preferences that we assume throughout the paper.

not, he consults the experts in order, stopping if one of them supports one of the r best projects. If, after round r, the manager has not received support for one of his r best projects, round $r + 1$ begins. The procedure stops after (at most) K rounds.

A special case provides intuition for the results. Suppose the projects are ordered so that the manager prefers the greatest project and experts have single-peaked preferences over projects; that is, expert i is characterized by an optimal project x_i^*. Her preferences increase for $x < x_i^*$ and decrease thereafter. In this setting, projects greater than x_i^* are weakly dominated and the salient prediction for the simultaneous-move game is that the manager will implement the maximum of the x_i^*. Even if preferences are not single peaked, the equilibrium cannot result in a project less than the maximum x_i^* because otherwise an expert would have a profitable deviation. Will the manager actually do better? The answer is plainly "yes" if the maximum x_i^* is not an equilibrium task. It will fail to be an equilibrium if there exists an expert j who prefers some project $x_j > \max_i x_i^*$ to $\max_i x_i^*$.

Both the simultaneous and sequential procedures that we consider limit the ability of the manager to restrict the projects that the experts can approve. Specifically, an expert can support any project when consulted. The manager would benefit from the ability to restrict recommendations. For example, suppose there are three projects: a status quo project, an intermediate project that is the favorite of all experts, and the manager's favorite project that exactly one expert prefers to the status quo. The intermediate project will be the equilibrium outcome that survives our refinement in the simultaneous-move game and also under the sequential protocol that we described. If experts could approve only the status quo or the manager's favorite outcome, however, then the manager would obtain his favorite. In Hu and Sobel (forthcoming), we permit the manager to prevent experts from approving certain projects. We show that limiting the options of the experts is strictly beneficial for the manager in both sequential and simultaneous consultations. We show that there is a sequential procedure with commitment that leads to the same project as the simultaneous-move game with commitment.

We organize the remainder of the paper as follows. Section I describes the basic model. Section II describes different interpretations of the model. Section III contains the analysis of the simultaneous-move game. Section IV contains the analysis of the sequential game. Section V describes related literature.

I. Underlying Strategic Environment

There is a finite set of players, who we call experts. I denotes the finite player set.[2] We assume that there is a finite set $X \subset \mathbb{R}$ available to each player. We call elements of X projects. We assume that X is ordered by the usual $\geq$ relation on $\mathbb{R}$.

For $\mathbf{x} = (x_1, \ldots, x_I)$,[3] $x_i \in X$ let $M(\mathbf{x}) = \max\{x_1, \ldots, x_I\}$. Each expert i has a payoff function $\tilde{u}_i : X^I \to \mathbb{R}$.[4] For each i, we assume that there exists $u_i : X \to \mathbb{R}$ such that $\tilde{u}_i : X^I \to \mathbb{R}$ is defined as $\tilde{u}_i(\mathbf{x}) \equiv u_i(M(\mathbf{x}))$. Throughout the paper, we

[2] In an abuse of notation, we also let I denote the cardinality of the player set.

[3] We use boldface to denote profiles that consist of actions or strategies of multiple players.

[4] Formally, the payoff function should be defined on strategy profiles. For the simultaneous game we study in Section III, X^I is the set of strategy profiles. When we study sequential games in Section IV, the strategy sets are more general, but we still denote payoff functions by $\tilde{u}_i$.

assume that the u_i are one-to-one. Because X is finite, u_i will be one-to-one for a generic set of preferences. (We say that a property is generic if it holds for an open set of Lebesgue measure one and the property is nongeneric otherwise. In this paper, genericity always refers to the property that utility functions are one-to-one. Hu and Sobel (forthcoming) analyze the model for more general X and nongeneric games.) We denote the minimum element of X by $\underline{x}$ and the maximum element by $\bar{x}$.

II. Interpretation of the Model

We study strategic interactions between the experts in the strategic environment described in Section I. The environment is abstract. We discuss several ways to interpret the environment in this section.

A. *Project Approval*

Assume that (in addition to the experts) there is a manager who strictly prefers larger projects. We interpret $\underline{x}$ as the status quo project. The manager prefers any other project to $\underline{x}$. The manager cannot implement a project different from the status quo without the assistance of at least one expert. We will study strategic environments in which experts announce which project they support. When offered a variety of projects, the manager will select the maximum (his most preferred project from the set). For this reason, we assume that experts report only a single project and that preferences over profiles $\mathbf{x} \in X^I$ depend only on the maximum component of $\mathbf{x}$. That is, we study a reduced form of a game in which the manager is a strategic player who selects his favorite project among those offered by experts. In this environment, there are no natural restrictions on the experts' preferences over X. For example, let $I = 2$ and $X = \{0, 0.1, \ldots, 0.9, 1\}$, where x describes a project that generates total surplus x. If the manager cares about total surplus, then he prefers x to x' if and only if $x > x'$. But different projects may distribute the share of the surplus across experts differently. This example suggests how transfers could be compatible with our framework as long as they are included in the description of elements of X.

There are settings in which the manager tries to get approval of projects simultaneously and others in which sequential consultation appears to be the rule.

B. *Asking for Permission*

One can interpret an expert's strategy as permission to undertake certain activities.[5] (If expert i supports a project, then the manager—who we think of as a decision-maker in this application—can pursue any project no better for the manager than this project.) Imagine that the decision-maker is a teenager and the experts are parents. The teenager requires a parent to give permission for an activity (the permission could be in the form of signing a waiver that allows the teenager to go on a school trip or permission to use a family car or stay out late). Alternatively, a manager may need

[5] We thank Inés Moreno de Barreda for this suggestion.

to secure necessary inputs from one of many divisions. The different divisions may be semiautonomous and have different preferences. In these settings, it is natural to assume that direct transfers are not feasible.

Our analysis identifies an equilibrium in which the decision-maker receives permission to do anything he wishes but also points out that this prediction is often implausible and identifies the equilibrium preferred by the experts as a more robust prediction. We think the sequential protocol introduced in Section IV provides an accurate description of what happens in settings like this.

C. *Agenda Manipulation*

There is a large literature on voting in committees (see, for example, Banks 1985 or Miller 1980).[6] In this literature, there are a finite set of projects (bills) and an odd number of committee members (experts). Experts have strict preferences over projects. An agenda is an ordering of the projects. An agenda induces a voting game in which the experts first choose via majority vote between the first and second elements in the agenda and continue so that in stage n, they choose between the winner of stage $n - 1$ vote and the $n + 1$ project. A project is a sophisticated outcome if it is the (refined) equilibrium associated with some agenda. The literature characterizes the set of sophisticated outcomes. The literature focuses on sequential procedures for a technical reason and a practical reason. The technical reason is that (when there is an odd number of voters) majority rule selects a unique winner of each pairwise contest but does not identify an outcome of simultaneous voting. This means that the simultaneous game is not well defined. The practical reason is that agendas are used in real legislatures.

One can view our problem as a version of the problem of voting on committees when it takes only one vote to advance an alternative. This change clearly makes it easier to obtain approval. It also provides an environment in which we can compare outcomes from simultaneous procedures to those from sequential procedures.

D. *Bayesian Persuasion*

We can interpret the model as a description of persuasion with many Senders. Assume that there is an underlying state of the world and experts provide the decision-maker with "experiments"—procedures that produce for each state of the world a probability distribution over a set of signals observable by the decision-maker. The decision-maker then makes a decision based on the signals he observes (and knowledge of the experiments and the prior distribution on the state of the world). This interpretation is consistent with the model of competition in persuasion in Gentzkow and Kamenica (2017b).

Let us describe the connection in somewhat more detail. We restrict attention to finite environments. In any Bayesian Persuasion problem, there is a given state space, Θ. We create a new state space $\Theta^* \equiv \Theta \times T$ where $(\theta, t) \in \Theta^*$, t is uniformly distributed on a finite set T, independent of θ. A partition of Θ^* is an experiment

[6] We thank referees for this suggestion.

(in the sense that observing an element of the partition generates a posterior distribution on Θ). Provided that we allow only finitely many experiments, the Bayesian Persuasion model is described by our model, although we must extend the analysis to partially ordered X.[7] In models of competition in persuasion, the decision-maker observes the realization of all experiments. In our one-dimensional model, observing the maximal experiment is equivalent to observing all experiments. In the multidimensional extension, when experiments need not be ordered, the common refinement of two partitions (the maximum) depends nontrivially on both partitions. Consequently, the ability to observe all experiments is necessary.

Gentzkow and Kamenica (2017a, 2017b) study a model in which experts simultaneously choose how much to communicate to a decision-maker in a Bayesian Persuasion framework. In these models, the decision-maker wants to know the value of the state of the world, and the strategies of experts are arbitrary signals (joint probability distributions on the state and message received by the decision-maker). Gentzkow and Kamenica (2017b) show that adding an agent may decrease the amount of information revelation but provide a condition under which increasing the number of experts increases the amount of information revealed. In our environment, additional experts are always valuable because Gentzkow and Kamenica's condition holds when experiments are completely ordered and all experts have access to the same set of experiments. Gentzkow and Kamenica do not focus on equilibrium selection, but they note the existence of multiple equilibria and the tendency of experts to prefer less disclosure. Li and Norman (2021) study a sequential version of the Gentzkow and Kamenica model. The paper provides an existence and partial characterization result. They show that sequential persuasion results in no more informative equilibria than simultaneous persuasion. Li and Norman (2018) also note that the order of disclosure matters, providing an example where inserting an additional expert into some (but not all) locations in a sequence may decrease the amount of information disclosure.

Ravindran and Cui (2022) study a Bayesian Persuasion problem in which Senders simultaneously select experiments. They show that if the preferences of game between the Senders is zero sum, then generically full disclosure is the unique equilibrium outcome. Ravindran and Cui note that the zero-sum property implies that all Senders are indifferent between any equilibrium payoff and the full disclosure payoff. Using Proposition 1, it is straightforward to show that this condition guarantees that the manager's favorite outcome is the unique outcome of the simultaneous-move game that survives IDWDS.

III. Simultaneous Moves

In this section, we study the game in which each expert simultaneously selects an element in X. If $\mathbf{x} = (x_1, \ldots, x_I)$ is the profile of projects, then expert i's payoff

[7] Hu and Sobel (forthcoming) study a model with a more general X. Assuming that experts' preferences are quasi supermodular, many of the basic insights from this paper continue to hold. It is straightforward to give conditions that guarantee quasi supermodularity in simple Bayesian Persuasion problems, but in general the condition is restrictive.

is $\tilde{u}_i(\mathbf{x}) = u_i(M(\mathbf{x}))$. We interpret the minimum element of X, $\underline{x}$, as a status quo. So when expert i wishes to support no project, she uses the strategy $x_i = \underline{x}$.

Section IIIA points out basic properties of the Nash equilibria of this game. Section IIIB describes the equilibrium refinement. Section IIIC states the characterization result.

A. *Basic Properties*

A profile $\mathbf{x}^* = (x_1^*, \ldots, x_I^*)$ with the property that $u_i(M(\mathbf{x}^*)) \geq u_i(M(x_i, \mathbf{x}_{-i}^*))$ for all x_i and all i is a Nash equilibrium profile. If $\mathbf{x}^*$ is a Nash equilibrium, we refer to $M(\mathbf{x}^*)$ as an equilibrium outcome. For any equilibrium profile $\mathbf{x}^*$, a strategy profile $\mathbf{x}$ that satisfies $x_i \leq x_i^*$ and at least two $x_j = M(\mathbf{x}^*)$ is a Nash equilibrium. The manager obtains his most preferred outcome when $M(\mathbf{x}^*)$ is equal to the maximum element in X, $\bar{x}$. Project $\bar{x}$ is always an equilibrium outcome. A strategy profile in which at least two experts play $\bar{x}$ supports it. Typically, there are other Nash equilibria.

We claim that the pure-strategy Nash equilibria are Pareto ranked from the perspective of the experts. In fact, if x^* and x^{**} are equilibrium outcomes and $x^{**} \geq x^*$, then all experts prefer x^* to x^{**}. To see this, observe that if any expert preferred the outcome x^{**} to x^*, then she could deviate by using the strategy x^{**} instead of the strategy she uses in the equilibrium that leads to the outcome x^*. Consequently, the equilibria are Pareto ranked. Furthermore, the manager's preferences are completely opposed to the (common) preferences of the experts.

B. *Weak Dominance*

The possibility of multiple equilibria leads us to consider a more restrictive solution concept.

DEFINITION 1: *Given subsets $X_i' \subset X$, with $X' = \prod_{i \in I} X_i'$, expert i's strategy $x_i^* \in X_i'$ is a best response to $\mathbf{x}_{-i} \in X_{-i}'$ relative to X_i if $\tilde{u}_i(x_i^*, \mathbf{x}_{-i}) \geq \tilde{u}_i(x_i, \mathbf{x}_{-i})$ for all $x_i \in X_i$. Expert i's strategy $x_i \in X_i'$ is weakly dominated relative to X' if there exists $x_i' \in X_i'$ such that $\tilde{u}_i(x_i, \mathbf{x}_{-i}) \leq \tilde{u}_i(x_i', \mathbf{x}_{-i})$ for all $\mathbf{x}_{-i} \in X_{-i}'$, with strict inequality for at least one $\mathbf{x}_{-i} \in X_{-i}'$.*

DEFINITION 2: *The set $S = S_1 \times \cdots \times S_I \subset X$ survives iterated deletion of weakly dominated strategies (IDWDS) if for $m = 0, 1, 2, \ldots$, there are sets $S^m = S_1^m \times \cdots \times S_I^m$, such that $S^0 = X$, $S^m \subset S^{m-1}$ for $m > 0$; S_i^m is obtained by (possibly) removing strategies in S_i^{m-1} that are weakly dominated relative to S^{m-1}; $S^m = S^{m-1}$ if and only if for each i no strategy in S_i^{m-1} is weakly dominated relative to S^{m-1}; and $S_i = \bigcap_{m=1}^{\infty} S_i^m$ for each i.*[8]

For finite games, it must be the case that there exists an m such that $S^r = S^m \neq \emptyset$ for all $r > m$. There are typically many different procedures that are consistent

[8] Our notation follows these rules: superscripts denote steps in an iterated process; subscripts denote players.

with Definition 2. These procedures may lead to different sets that survive the process, but Proposition 1 shows under our assumptions all sets that survive lead to the same maximum project.

We analyze the implications of applying iterated deletion of weakly dominated strategies. Sobel (2019) introduces a class of games called WID-supermodular games and describes general properties of strategies that survive the process of iteratively deleting weakly dominated strategies in these games. He shows that the simultaneous-move game is a WID-supermodular game and provides the characterization result that we describe next.

C. *Characterization*

This section characterizes the unique outcome that survives IDWDS in the generic simultaneous-move game. We begin with some general properties of the equilibrium set.

We have observed that the maximal project is an equilibrium outcome. There must be a minimum equilibrium outcome because X is completely ordered and finite. We next show that the manager prefers every project that survives iterated deletion of weakly dominated strategies, whether it is an equilibrium outcome or not, to the minimum equilibrium outcome. Before we describe the result, we let

$$\pi^* = \min\left\{\pi \in X : u_i(\pi) \geq u_i(x) \text{ for all } x > \pi \text{ and all } i\right\}.$$

The outcome π^* is Pareto efficient (from the perspective of the experts) in the set of Nash equilibrium payoffs.

We note several consequences of this definition. It is immediate that if π is an equilibrium outcome, then $\pi \geq \pi^*$. Furthermore, if $\bar{x} > \pi^*$, then there will be equilibrium outcomes greater than π^*. In fact, if π is an outcome such that for all i, $u_i(\pi') \leq u_i(\pi)$ for all $\pi' \geq \pi$, then any strategy profile $\mathbf{x}$ that satisfies $x_i \leq \pi$ and at least two $x_j = \pi$ is a Nash equilibrium. In particular, $\bar{x}$ is always an equilibrium outcome. The next result asserts that outcomes $\pi > \pi^*$ fail to satisfy a refinement.

PROPOSITION 1: *If* $\mathbf{x}$ *is a strategy profile that survives IDWDS in the simultaneous-move game, then* $M(\mathbf{x}) = \pi^*$.

The proposition identifies a unique outcome that survives iterated deletion of weakly dominated strategies. Sobel (2019) contains a proof of the proposition. Hu and Sobel (forthcoming) extend the result to nongeneric preferences and incompletely ordered X.

IV. Sequential Protocols

The manager's preferred outcome does not always survive iterated deletion of weakly dominated strategies when experts move simultaneously. This leaves open the question of whether the manager could do better by consulting the experts in a different way. This section discusses the issue. We introduce a family of sequential protocols and describe a simple member of the family that performs at least as well as any other sequential protocol (from the standpoint of the manager). This protocol generates the same outcome as the simultaneous-move game.

The online Appendix contains an example that illustrates that simple sequential procedures may not lead to good outcomes for the manager. Section IVA describes general sequential procedures and introduces the canonical procedure. Section IVB shows that from the perspective of the manager, the canonical procedure performs at least as well as any other sequential procedure and as well as simultaneous consultation. Section IVC discusses some comparative-statics properties.

A. *Definition of Sequential Protocols*

In a sequential protocol, the manager selects an expert; the expert can then decide to approve any project in X; the manager then decides whether to implement a project that has been approved or to move to another expert. We assume that protocols are deterministic; that they must end after a finite number of consultations; and that when the consultation procedure ends, the manager implements the largest project that has been approved. The restriction to deterministic protocols simplifies the exposition. In our model, no stochastic protocol can do strictly better than the deterministic protocol we describe. We do not know if potentially infinite protocols can benefit the manager. The restriction that the manager implements the largest project that has been approved parallels the assumption we made for the simultaneous game. We focus on a particular protocol, which we call the canonical protocol (CP).

Suppose the projects can be ranked in the order of the manager's preferences: $\bar{x} = \pi_K \succ \pi_{K-1} \succ \cdots \succ \pi_1 = \underline{x}$.

DEFINITION 3: *The canonical sequential protocol has at most K rounds, starting with round $r = 1$. In round r, the manager consults experts in order. If any expert in round r chooses a project $\pi \succeq \pi_{K-r+1}$, the process stops and the manager implements the largest project supported. Otherwise, round $r + 1$ begins.*

We describe protocols and CP formally in the online Appendix.[9]

The ability to create a protocol assumes that the manager has commitment power, but the commitment power is limited. A protocol specifies rules for consultation. The rules specify who the manager consults and when he terminates consultation. This commitment power may be valuable to the manager. For example, suppose $K = 3$, $u_1(\pi_2) > u_1(\pi_1) > u_1(\pi_3)$, and $u_2(\pi_1) > u_2(\pi_3) > u_2(\pi_2)$. In the second round, expert 1 knows that if she approves π_2, then the manager would like to consult expert 2 again because expert 2 prefers π_3 to π_2. But if expert 1 refuses to approve π_2, the protocol specifies that the manager move to expert 2, who would prefer to wait until round 3 than support π_2. Hence, if the manager lacked commitment power, expert 2 would not support π_3 in the first round. We believe that it is realistic to assume that managers have the power to discontinue consultations but not the power to demand approval of a particular project.[10]

[9] We are grateful to Christopher Turansick for suggesting this procedure.

[10] It is common to have rules governing consultation procedures. Robert's Rules of Order (De Vries 1998), which establishes rules governing who can speak and what can be discussed in a meeting, is a leading example.

We limit commitment in two ways. First, the project the manager takes must be the best available given the strategy of the experts. By making this assumption, we implicitly assume that the manager cannot commit to implementing a project that he likes less than a project he could implement. Second, we assume that the manager cannot restrict the set of projects that an expert can support. An extreme way to do this would be to exclude some elements of X from every choice set. The manager cannot restrict the experts' strategies and must take the best available project in the simultaneous game, where experts can support any project and the manager's action is the maximum project supported. Hence, the limits to commitment make the sequential game comparable to the simultaneous game. These restrictions on commitment ability are in the spirit of restrictions that sequential rationality would impose if we model the manager as a strategic player.[11] Hu and Sobel (forthcoming) show that the ability to make commitments is valuable.

B. *Performance of Sequential Protocols*

In this section, we show that the CP generates π^*, which is the outcome that survives IDWDS in the simultaneous-move game and that no sequential protocol does better. Hence, simultaneous consultation performs as well as, but no better than, a well-designed sequential protocol.

PROPOSITION 2: *If π is a project that survives IDWDS in the game determined by CP, then $\pi = \pi^*$.*

Proposition 2 identifies the manager's outcome for the canonical sequential protocol. For every $\pi < \pi^*$, there will be an expert who strictly prefers a higher project. From this observation, it is straightforward to show that a project π such that $\pi < \pi^*$ cannot be generated by CP. By the definition of π^*, no higher outcome is possible.

PROOF:

Let π be a project that survives IDWDS. Suppose it is generated by the strategy profile $\hat{\mathbf{s}}$. Suppose $\pi < \pi^*$. We will show that there exists an expert i such that $\hat{s}_i$ is weakly dominated.

By the definition of π^*, it must be the case that for some i,

$$\text{there exists } x_i \text{ with } x_i > \pi \text{ such that } u_i(x_i) > u_i(\pi). \tag{1}$$

Find a history $\hat{h}$ consistent with $\hat{\mathbf{s}}$ such that the manager consults expert i at history $\hat{h}$ and if i's play at $\hat{h}$ is $\hat{s}_i(\hat{h})$, then there is no undominated strategy profile that consults i again. It is possible to find such a history because $\pi < \pi^*$ implies that the manager must consult every expert at least once and because the protocol never consults an expert more than K times.

[11] The manager is not a strategic player, so we do not have a formal result that justifies our (lack of) commitment assumption as a reduced-form equilibrium of a model with a strategic manager.

Consider an alternative strategy of expert i in which

$$s_i'(h) = \begin{cases} x_i' & \text{if } h = \hat{h} \\ \hat{s}_i(h) & \text{otherwise}' \end{cases}$$

where x_i' solves $\max u_i(x_i)$ subject to $x_i > \pi$. We know that x_i' exists and satisfies $u_i(x_i') > u_i(\pi)$ by (1). By the definition of $\hat{h}$, if expert i supports x_i' after $\hat{h}$, the protocol must stop. (We know that the protocol would stop in the next round with the outcome π, so it must stop immediately when expert i supports something strictly better for the manager than π.) It follows that s_i' weakly dominates $\hat{s}_i$. The strategy s_i' does exactly as well as $\hat{s}_i$ for any strategy profile that does not induce the history $\hat{h}$. We know that some strategy profile does induce $\hat{h}$ and, by construction, expert i does strictly better in any such case. Consequently, any outcome $\pi < \pi^*$ must be generated by a strategy profile in which one player uses a weakly dominated strategy. This establishes that CP generates an outcome π that satisfies $\pi \geq \pi^*$. The proposition follows from Proposition 3. ∎

The next result states that π^* is an upper bound for all sequential protocols because the unique subgame-perfect equilibrium determines the unique outcome that survives IDWDS in perfect-information games.

PROPOSITION 3: *If* $\pi > \pi^*$, *then there exists no sequential protocol that generates the project* π *in a pure-strategy, subgame-perfect equilibrium.*

PROOF:

Proposition 3 follows from backward induction. After each history that supports no more than π^*, it is never a best response to approve more than π^*. So if the first expert anticipates that the final project supported will be more than π^*, then she can do strictly better by approving π^* and no one else will add more to π^*. Consequently, there will never be projects greater than π^* in equilibrium. ∎

C. *Comparative Statics*

In this section, we make a few observations about the value of adding experts.

Adding an expert cannot harm the manager in the sense that if π is a project that survives IDWDS for the original set of experts in the simultaneous-move game, a project at least as good as π for the manager will survive if additional players are added; the new player need not be consulted in a sequential protocol. In a model of Bayesian Persuasion closely related to our model (see Section IID for a comparison), Li and Norman (2018) show that adding an expert may hurt the decision-maker if the expert must be inserted in a particular place.

Adding an additional expert is beneficial if and only if doing so increases π^*. An expert who does not increase this quantity is **redundant**. If preferences are single peaked, all experts except the one with the greatest peak is redundant. More generally, if there is a pair of experts i and j such that for all $x' \succ x$, $u_i(x') > u_i(x)$ whenever $u_j(x') \geq u_j(x)$, then expert j is redundant.

V. Related Literature

We know of several papers that compare simultaneous to sequential interactions in different contexts. Dekel and Piccione (2000) compare simultaneous to sequential voting institutions. There are two options and a finite number of voters. Voters can either vote for or against the status quo. Voters do not know their valuations but receive private signals. The authors compare the equilibria of games in which voters cast votes simultaneously to those in which votes are sequential. They show that a symmetric informative equilibrium of the simultaneous game is an equilibrium to any sequential game. Weaker results hold for asymmetric equilibria.[12] Although that paper reaches a conclusion that is similar to ours, we do not see a formal connection between the analyses. The model of Dekel and Piccione (2000) focuses on the possibility of learning something about the state from the behavior of other voters. Our experts lack private information. Our equivalence result requires an equilibrium refinement and commitment power in the design of sequential mechanisms. Schummer and Velez (2021) identify conditions under which social choice functions that can be implemented in truthful strategies when players move simultaneously cannot be truthfully implemented when players move sequentially. The context is quite different from our paper, but it suggests environments in which sequential procedures will perform less well than simultaneous ones.

Doval and Ely (2020) and Salcedo (2017) characterize all equilibria that can arise from some information structure and some extensive form (for a fixed set of players and preferences over final outcomes). Their construction involves a "canonical extensive form" that is sufficient to generate any equilibrium. In a canonical extensive form, each player moves at most once. Our construction requires that an individual player may move more than once. The reason for this difference is that Doval and Ely's construction requires a partial commitment assumption that requires that once a player has made an action choice, that player can have no other payoff relevant moves. This assumption does not hold in our model.

Armstrong and Vickers (2010) study a delegation problem with a single agent. They assume that there is a set of potential projects. The principal selects a set of permitted projects. Nature then determines that set of potential projects that are actually feasible. The agent then selects a project from the set of projects that are both permitted and feasible (or does not select any project). In the basic model, the principal learns the characteristics of the selected project but does not learn which projects are feasible. Armstrong and Vickers describe the solution to this problem, which typically involves restrictions on the set of permitted projects. Our focus is on how the decision-maker can leverage differences in preferences between experts to improve his outcome, while Armstrong and Vickers assume that there is a single agent. A common feature of their approach and our analysis, especially in the case where commitment is feasible, is that the principal gains from having the option to limit the choice of the agent or experts.

Goel and Hann-Caruthers (2020); Guo and Shmaya (2021); and Kartik, Kleiner, and Van Weelden (2021) study mechanism design problems without transfers under

[12] Dekel and Piccione (2014) study a voting model in which the timing of votes is a strategic choice.

incomplete information in which a decision-maker must get the approval of a single agent to carry out a project. These papers differ from ours because they give a player much greater commitment power, assume incomplete information, and do not discuss the impact of having multiple experts.

REFERENCES

Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica* 78 (1): 213–44.

Banks, J.S. 1985. "Sophisticated Voting Outcomes and Agenda Control." *Social Choice and Welfare* 1 (4): 295–306.

Dekel, Eddie, and Michele Piccione. 2014. "The Strategic Dis/advantage of Voting Early." *American Economic Journal: Microeconomics* 6 (4): 162–79.

Dekel, Eddie, and Michele Piccione. 2000. "Sequential Voting Procedures in Symmetric Binary Elections." *Journal of Political Economy* 108 (1): 34–55.

De Vries, Mary A. 1998. *The New Robert's Rules of Order*. New York: Signet.

Doval, Laura, and Jeffrey C. Ely. 2020. "Sequential Information Design." *Econometrica* 88 (6): 2575–608.

Gentzkow, Matthew, and Emir Kamenica. 2017a. "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces." *Games and Economic Behavior* 104: 411–29.

Gentzkow, Matthew, and Emir Kamenica. 2017b. "Competition in Persuasion." *Review of Economic Studies* 84 (1): 300–22.

Goel, Sumit, and Wade Hann-Caruthers. 2020. "Project Selection with Partially Verifiable Information." arXiv: 2007.00907.

Guo, Yingni, and Eran Shmaya. 2021. "Project Choice from a Verifiable Proposal." Unpublished.

Hu, Peicong, and Joel Sobel. Forthcoming. "Getting Permission When Options Are Partially Ordered." In *David Gale: Mathematical Economist: Essays in Appreciation on His 100th Birthday*, edited by M. Ali Khan, Nobusumi Sagara, and Alexander J. Zaslavski. Cham, Switzerland: Springer International Publishing.

Kartik, Navin, Andreas Kleiner, and Richard Van Weelden. 2021. "Delegation in Veto Bargaining." *American Economic Review* 111 (12): 4046–87.

Li, Fei, and Peter Norman. 2018. "On Bayesian Persuasion with Multiple Senders." *Economics Letters* 170: 66–70.

Li, Fei, and Peter Norman. 2021. "Sequential Persuasion." *Theoretical Economics* 16 (2): 639–75.

Miller, Nicholas R. 1980. "A New Solution Set for Tournaments and Majority Voting: Further Graph-Theoretical Approaches to the Theory of Voting." *American Journal of Political Science* 24 (1): 68–96.

Ravindran, Dilip, and Zhihan Cui. 2022. "Competing Persuaders in Zero-Sum Games." Unpublished.

Salcedo, Bruno. 2017. "Interdependent Choices." Unpublished.

Schummer, James, and Rodrigo A. Velez. 2021. "Sequential Preference Revelation in Incomplete Information Settings." *American Economic Journal: Microeconomics* 13 (1): 116–47.

Sobel, Joel. 2019. "Iterated Weak Dominance and Interval-Dominance Supermodular Games." *Theoretical Economics* 14 (1): 71–102.