Submission Number: PET11-11-00220

Relying on Experts with Simple Contracts

Alper Nakkas
Vanderbilt University

Abstract

Decision makers often rely on recommendations of experts, who are more efficient and productive at processing noisy information. This paper investigates the consequences of delegating decision power to an expert in return of a simple contract when the expert has an ability to create a surplus at each state of the uncertainty. We identify that surplus bringing behavior together with asymmetric information leads to two different types of effects, which we refer to as adjustment effect and exclusion effect. While the direction of the adjustment effect may be ambiguous, the exclusion effect causes a divergence from the optimal decision. We show that delegation of decision power can, in some circumstances, benefit not only the decision maker and the expert but also other parties that may be affected by the decision.

Submitted: March 12, 2011.

Relying on Experts with Simple Contracts*

Alper Nakkas†
Vanderbilt University

Abstract

Decision makers often rely on recommendations of experts, who are more efficient and productive at processing noisy information. This paper investigates the consequences of delegating decision power to an expert in return of a simple contract when the expert has an ability to create a surplus at each state of the uncertainty. We identify that surplus bringing behavior together with asymmetric information leads to two different types of effects, which we refer to as adjustment effect and exclusion effect. While the direction of the adjustment effect may be ambiguous, the exclusion effect causes a divergence from the optimal decision. We show that delegation of decision power can, in some circumstances, benefit not only the decision maker and the expert but also other parties that may be affected by the decision.

1 Introduction

Decision under uncertainty becomes more and more difficult as the amount of information needs to be processed in order to reach reliable conclusions expands. As a natural result, decision makers often rely on recommendations of experts, who are more efficient and productive at processing noisy information. While an expert can be thought as a profit maximizing entity, in many economic circumstances, it may simply be a person or firm that benefits from the better judgements of the decision maker.

There are many situation in which the decision makers rely on the experts, even without paying any virtual costs. Individuals almost always seek advices from their family members, friends, and colleagues, especially for big decisions that involve uncertainties. In most trials, litigants delegate their decision powers concerning their cases to lawyers because the lawyers, as experts, not only improve the likelihood of winning the clients' cases, but also have a positive impact on the amount of the award that the plaintiff will receive in the case of winning the trial. Doctors, mechanics..etc are other good examples where the decision maker's rely on the experts' opinions.

*This is a very preliminary draft. Please do not circulate.

†Owen Graduate School of Management, Nashville, TN 37203, USA. Contact: alper.nakkas@owen.vanderbilt.edu.

Firms also rely on the recommendations of the experts in many situations. For example, in the retailing industry, many firms seek strategic recommendations from their leading manufacturers regarding key category management decisions such as category assortment, a practice often referred as category captainship. The category captains are not paid employers of the retailers but they benefit indirectly from any increase in the category demand. For example, Coors Brewing, which is the third-largest brewing company in the United States, helps for a number of its retail clients in the alcoholic beverages category to improve their profitability. The key insight provided by Coors Brewing company to one of their retail clients was that the retail chain's core shopper best matched the characteristics of premium light beer purchasers. However, the retailer was not able to convert its shoppers into premium light beer buyers. In addition, they did a study on what products perform well in a particular region and developed best practice planograms for the retailers for which they served as a category captain (Progressive Grocer 2007). For another retailer, Small Planet strategists found that placement of the organic/natural items was not a critical success factor when compared to things such as the type of consumer (heavy vs light), variety, and duration of shopping trip (Progressive Grocer 2007).

When there are uncertainties associated with the decision making process, an expert can be helpful in different dimensions such as generating key insights to improve operation aspects or surplus enhancing strategies. In this paper, we focus on two types of benefits that the experts can provide to the decision makers. The experts usually (1) provide insights which are not readily available to the decision makers and help decision makers take a more informed action and/or (2) improve the inefficiencies in the decision making process, which corresponds to an increase in the decision maker's surplus. The key consumer insights provided by Coors Brewing is an example of former type of benefit, whereas a lawyers ability to determine the evidences that improve the likelihood of winning the trial is an example of the latter type of benefit.

Ideally, a decision maker who plans to delegate her decision making privileges to an expert (or, equivalently, follow the expert's recommendations) would write a contract that aligns the incentives of the expert. There is a vast literature that focuses on this issue. However, it is not always possible for decision makers to write optimal contracts. In many cases, decision makers prefer simple contracts which are easy to implement and monitor. Motivated by the existence of imperfect contracts, the goal of this paper is to better understand the consequences of experts for decisions under uncertainty. To this end, we consider a model where a decision maker relies on an expert for a decision under uncertainty. First, we consider a model where the decision maker is responsible for taking the optimal action when the decision is surrounded by uncertainty. Then, we consider a simple delegation game where the decision maker relies on the decision of the expert. We assume that the decision maker benefits from using an expert for decisions because the expert can (1) provide insights not readily available to the decision maker and (2) increase the surplus. Our results are based on a comparison of these two models. We show that surplus bringing behavior together with asymmetric information leads to two different types of effects, which we refer *adjustment effect* and *exclusion effect*. While the direction of the adjustment effect may be ambiguous for the decision, the exclusion effect causes a divergence from the optimal decision.

2 The Model

A decision maker (D) faces an optimization problem, which inherently depends on a random event m. The decision maker observes the value of the random event, m, with density $f(m)$. The differentiable cumulative distribution function, $F(m)$, is supported on $[0,1]$. Let μ and σ be the mean and variance of m, respectively. The decision maker has a twice continuously differentiable von Neumann-Morgenstern utility function $U(y,m,n)$, where $y \in \mathcal{R}$, a real number, is the action taken by the decision maker and $n \in \mathcal{R}$ measures the information processing ability of the decision maker. Throughout the paper we shall assume that, for each m, denoting partial derivatives by subscripts in the usual way, $U_1(y,m,n) = 0$ for some y, and $U_{11}(\cdot) < 0$, so that U has a unique maximum for each given m and n; and that $U_{12}(\cdot) < 0$ and $U_{13}(\cdot) > 0$. The latter conditions ensure that the best value of y from a fully informed agent's standpoint is a strictly decreasing function of the true value of m and strictly increasing function of n, respectively.[1] For simplicity, we assume that U is linear in the random variable m. In order to differentiate between directions of the actions, we say that the action y' is *softer* than y if $y' < y$ and *more aggressive* if $y' > y$.

2.1 Decision under Uncertainty

We start with the decision maker's problem when she chooses the optimal action under uncertainty. The decision maker maximizes her expected utility by solving

$$\max_y \int_0^1 U(y,m,n) f(m) dm. \tag{1}$$

Let y_D be the decision maker's optimal choice of action when the decision maker chooses her action under uncertainty.

Proposition 1. *The decision maker's optimal action under uncertainty, $y_D = y(\mu, n)$, is a decreasing function of unconditional expected event $\mu = E[m]$ and increasing function of the decision maker's information processing ability n.*

[1]These assumptions are not critical for the purpose of the paper. We could have similar results with different assumptions on cross partial derivatives. We use these assumptions only to have a clear exposition.

In order to choose the optimal decision, the decision maker averages all possible marginal benefits and costs in all states of the random variable m and balances these averages. This averaging behavior of the decision maker naturally leads to a sub-optimal choice of action. Due to the sub-optimality of the action to be taken, there is room for an improvement in the outcome and the decision maker has incentives to seek help from an expert. The decision maker's optimal action under uncertainty, y_D, is determined by the decision maker's ability to process information and the moments of the distribution of random variable m.

2.2 Decision with an Expert

We now consider the situation that the decision maker seeks help from an expert (E). We capture the expertise of the expert by assuming that the expert has private information about m, whereas the decision maker observes the value of a random variable, m, with density $f(m)$. The decision maker's utility $U(y, m, \delta n)$ is determined by the action taken by the decision maker upon receiving the expert's recommendation, the true value of m, and the expert's information processing ability δn. We assume that $\delta \geq 1$ in order to capture the expert's superior efficiency in processing the information. Therefore, the expert is beneficial for the decision maker not only because he eliminates the randomness from the decision process but also he generates a surplus for the decision maker.

We focus on the extreme situation that the expert's utility enhancing ability is known to the decision maker and it is presumably a result of the cooperation between the expert and the decision maker.[2] For instance, a category captain increases demand...etc. In practice, the expert's recommendation can be implemented immediately or can be censored by the decision maker (e.g., getting a second opinion). For simplicity, we assume that it is too costly for the decision maker to censor the recommendations of the expert. The expert's utility is presumably a function of the action he recommends as an expert, the true value of the state, and his information processing ability and denoted by $V(y, m, \delta n)$. We assume that $V(y, m, \delta n)$ is continuous in y and, for each m and δn, a unique $\bar{y}_E = \arg\max_y V(y, m, \delta n)$ exists.

The decision maker and the expert plays a delegation game. The sequence of events in the game is as follows. First, the expert observes the true state of m (i.e., his "type"). Then, the decision maker offers a *simple* contract. We define a contract as simple if the contract requires a single threshold. Thresholds here are considered on a broad spectrum and can be imposed on variables such as profit, volume of sales, number of products, quality level...etc. The utility that is generated by the threshold is denoted as T. After a utility threshold T is offered, the expert either accepts or rejects the offer. If the offer is accepted, then the expert recommends an action y, which in return is implemented by the decision maker. For simplicity, we assume that if the expert accepts

[2]We consider a situation in which the expert's ability is uncertain for the decision maker in the robustness section.

the contract and cannot achieve the goal set by the decision maker, then the expert pays a very high penalty to the decision maker. If the expert rejects the contract offer, then the decision maker updates her beliefs about m and decides on the action y. Technically, we model the delegation game as a screening game in which the (uninformed) decision maker makes a take-it-or-leave-it offer to the (informed) expert. All aspects of the game except m are common knowledge. We are interested in the pure strategy perfect Bayesian equilibria of the game.

2.3 The First Best Benchmark

Before we start the analysis of the delegation game, it is useful the determine the first best benchmark for the decision maker. The first best outcome of the decision maker is a solution of the decision maker's problem when she knows the true state of m and, moreover, she is as efficient as the expert in processing information. Thus, the first best benchmark is the solution of the problem

$$\max_{y} U(y, m, \delta n). \tag{2}$$

Let y_F be the optimal action in the first best benchmark. Because there is no uncertainty, the decision maker does not average the marginal benefits and costs across all states of the random event but rather balances marginal benefit and cost at the true state m while enjoying the increased information processing ability.

Proposition 2. *The decision maker's first best action, $y_F = y(m, \delta n)$, is a decreasing function of the event m and increasing function of the expert's information processing ability δn.*

2.4 Delegation Game

In order to determine whether the expert accepts a given utility threshold T, we first consider the optimal decision of the expert. The expert's problem for a given level of the threshold T is

$$\begin{aligned} \max_{y} \quad & V(y, m, \delta n) \\ \text{s.t.} \quad & U(y, m, \delta n) \geq T. \end{aligned}$$

The expert chooses his optimal action from the feasible set that is determined by the expert's ability to generate extra surplus as well as the target determined by the decision maker. The following proposition shows that the expert has a unique best response for each threshold T set by the decision maker.

Lemma 1. *There exists a unique solution $y_E = y(m, \delta n, T)$ that is monotonic in T for the expert's problem.*

Intuitively, the expert chooses his optimal action under the constraint that the decision maker is satisfied with the outcome of the action. Let $Y(m, \delta n, T)$ be the set of feasible actions for the expert. Notice that because U is a concave function, $Y(m, \delta n, T)$ is a closed range on the real line. Let y_E^L and y_E^H be the lower and upper bounds of $Y(m, \delta n, T)$, respectively. If the expert's utility is monotonic on $Y(m, \delta n, T)$, then the optimal choice of the expert is on the boundary and either y_E^L or y_E^H. However, if the expert's utility is not monotonic on $Y(m, \delta n, T)$, then the expert chooses an action in the interior of $Y(m, \delta n, T)$ since V is a continuous function and Y is a closed set. The decision maker is better off in the latter situation, simply because the incentives of the decision maker and the expert are better aligned. Then, the optimal action that the expert chooses has the following form:

$$y_E = y(m, \delta n, T) = \begin{cases} y_E^L, & \bar{y}_E < y_E^L; \\ \bar{y}_E, & \bar{y}_E \in [y_E^L, y_E^H]; \\ y_E^H, & \bar{y}_E > y_E^H. \end{cases} \tag{3}$$

The function y_E represents the best response behavior of the expert for a given threshold under the condition that the expert has already accepted the decision maker's offer. Notice that y_E is increasing in T when $\bar{y}_E < y_F$ since y_E^L is increasing in T. Similarly, y_E is decreasing in T when $\bar{y}_E > y_F$ since y_E^H is decreasing in T. Intuitively, by increasing the threshold, the decision maker shrinks the set of feasible actions for the expert, which forces the expert to recommend an action that is closer to the decision maker's ideal action. Throughout the rest of the paper, without loss of generality, we assume that $\bar{y}_E < y_F$ for all types of the expert.

Anticipating the best response behavior of the expert, the decision maker tailors her target selection. The decision maker screens the types of the expert by choosing a *borderline type*. Let m_b be the borderline type of the expert, who can deliver at most T. That is, for a given threshold T, the borderline type m_b is the value of m that satisfies $U(m, \delta n) = T$. Thus, the range $[0, m_b]$ is the set of expert types that cannot accept the utility threshold T when the expert's ability is δn. The borderline type chooses an action that coincides with the decision maker's first best action when the true state is m_b and the expert's information processing ability is δn. Notice that the decision maker does not necessarily try to choose an action that would leave the expert indifferent between accepting or rejecting the contract, which is typical in principal agent models. The decision maker rather chooses an action that brings the maximum possible surplus when the expert's incentive compatibility constraint is never binding. Let $T(m_b, \delta n)$ be the utility threshold that generates m_b as the borderline type when the expert's information processing ability is δn.

Lemma 2. *$T(m_b, \delta n)$ is an increasing function of both m_b and δ.*

Lemma 2 shows that there is a one-to-one correspondence between m_b and $T(m_b, \delta n)$ for any

$\delta \geq 1$. So, the decision maker's problem can be written as choosing her optimal action by solving the problem

$$\max_{m_b} \quad \int_0^{m_b} U(\mu_b, n) f(m) dm + \int_{m_b}^1 T(m_b, \delta n) f(m) dm \tag{4}$$

where $\mu_b = E[m|m \leq m_b]$ is the expected value of m when the decision maker knows that the random event m is in the range of $[0, m_b]$. The first term of the decision maker's expected profit is the decision maker's utility when the expert rejects the offer and the second term is her utility when the expert accepts the offer. While in the latter case the decision maker simply receives the utility she asks for, in the former case she updates her beliefs about the true state of m and maximizes her expected utility over the updated set of states. The reason for the change in the decision maker's optimal action in case of a rejection is that the decision maker updates her beliefs about m as a Bayesian agent. In particular, once the decision maker offers a contract that is characterized by m_b and the expert rejects it, the decision maker would infer that the true state of m cannot be higher than m_b since the expert would certainly accept the offer in such a case.

Lemma 3. *The decision maker's optimal strategy is to choose the borderline type m_b^* that solves the first order condition*

$$\frac{\partial U(m_b, \delta n)}{\partial m_b} = \lambda(m_b) \left[U(m_b, \delta n) - U(\mu_b, n) - m_b f(m_b) \frac{\partial U(\mu_b, n)}{\partial \mu_b} \right]$$

where $\lambda(m) = \frac{f(m)}{1-F(m)}$ is the hazard rate *function associated with the distribution of m.*

The decision maker chooses the utility threshold that balances the total marginal cost and total marginal benefit. On one hand, when the decision maker increases the borderline type marginally, she benefits from (1) the decrease in the noise related to the true state of m in case of a rejection and (2) the increase in the utility she will receive in case of an acceptance. While in the latter case the increase in the utility is simply because at each agreement the decision maker will receive more utility, in the former case it is due to the increased expected utility when the decision maker decides on the action without seeking any help from the the expert. On the other hand, an increase in the borderline type leads to a loss for the decision maker due to the possibility that the expert rejects the contract for additional types. For these types, the decision maker loses the possible improvement that the expert could have brought. Let m_b^* be the optimal borderline type and $T^* = T(m_b^*, \delta n)$ be the utility threshold that generates m_b^* as the borderline type.

3 Equilibrium and Results

We start with the definition of pure strategy perfect Bayesian equilibrium in our context. A strategy profile is defined as $S = (S_D, S_E)$ where S_D and S_E represent the strategies of the decision maker and the expert, respectively. The decision maker's strategy S_D is consist of her utility threshold T choice and the action choice when the expert rejects the utility threshold. The expert's strategy S_E is consist of a decision of accept (A) or reject (R) for each type of the expert represented by $\phi(m)$ and the action taken by each type when the utility threshold is accepted. Let $U(S)$ and $V(S)$ be the utilities that the decision maker and the experts receive, respectively, when the strategy profile is S. The beliefs of the decision maker about the expert's type when the expert rejects the offer is represented by $\xi(m)$, which attains a probability to each type m.

Definition 1. *A strategy profile $S^* = (S_D^*, S_E^*)$ and the belief structure ξ^* is a perfect Bayesian equilibrium if and only if it satisfies the following conditions:*

Sequential Rationality: Given ξ^*,

$$\begin{aligned} E[U(S_D^*, S_E^*)] &\geq E[U(S_D, S_E^*)] \quad \textit{for any} \quad S_D \neq S_D^* \\ E[U(S_D^*, S_E^*)] &\geq E[U(S_D^*, S_E)] \quad \textit{for any} \quad S_E \neq S_E^* \end{aligned}$$

Bayes Consistency of Beliefs: Given S^*, the decision maker's belief system $\xi^*(m)$ is consistent wherever possible. That is, $\int_0^1 \xi^*(m)dm = 1$ at each information set of the decision maker.

The above definition embeds the notion of the perfect Bayesian equilibrium in the following sense. First, the decision makers strategy is a best response to what she knows at that point, what the expert optimizes, and to her beliefs on the actual state of the random event. Second, the experts strategy is a best response to what he knows at that stage (the decision maker's utility threshold request and true state of the random event) and his conjectures on the beliefs of the decision maker. Finally, the decision makers actual beliefs and the experts conjectures on the decision makers beliefs coincide and the probability assigned to every node is computed as the probability of that node being reached given the strategy profile, i.e., Bayes' rule.

We can now summarize the equilibrium behaviors and payoffs of the decision maker and the expert. The following proposition shows that the strategies that are derived in the previous section are indeed form a perfect Bayesian equilibrium with a reasonable belief structure.

Proposition 3. *Let $S_D^* = (T^*, y(\mu_b^*, n))$, where $\mu_b^* = E[m|m \leq m_b^*]$, $S_E^* = (\phi^*(m), y(m, \delta n, T^*))$, and $\phi^*(m) = A$ when $m \geq m_b^*$ and $\phi^*(m) = R$ otherwise. Then, the strategy profile $S^* = (S_D^*, S_E^*)$ together with $\xi^*(m)$ is a perfect Bayesian equilibrium, where $\xi^*(m) = f(m)$ at the initial information* *set of the decision maker and*

$$\xi^*(m) = \begin{cases} \frac{f(m)}{F(m_b^*)}, & m \leq m_b^*; \\ 0, & m > m_b^*, \end{cases}$$

when the expert rejects the utility threshold offer T^*, i.e., $\phi^*(m) = R$.

In the equilibrium, the decision maker offers the utility threshold T^* and the expert accepts the offer if $m \geq m_b^*$ and rejects it otherwise. By setting a utility threshold, the decision maker benefits from two channels. First, even if the expert rejects the decision maker's offer, the decision maker learns something about the random event and makes a better informed decision in cases where the expert rejects the offer. Second, in cases where the expert accepts the offer, the decision maker enjoys a higher utility than the benchmark utility. We refer the utility gains of the decision maker as a result of the first effect as the information value of the delegation game, whereas the second one as the utility value of the game. Lemma 2 shows that the information value of the delegation game is greater than the utility value for relatively low values of δ. This is because, as δ increases, the expert's help becomes more attractive, which increases the relative importance of the information value against the utility value. Thus, the decision maker has to make a tradeoff between the information and utility values of the delegation game.

Our next result is about the impact of the delegation of the decision power on the action taken. Because the incentives of the decision maker and the expert are not perfectly aligned, the expert tries to take an action that is closest to his ideal action, whereas the decision maker tries to make the expert to take the action that is closest to her ideal action by setting a suitable target. The following proposition summarizes the difference between the actions taken by the decision maker and the expert. Let y_E^* be the resulting action of the delegation game.

Proposition 4. *There exists a type m^* such that (i) if $m \in [0, m_b^*)$, then $y_E^* > y_D$, (ii) if $m \in [m_b^*, m^*]$, then $y_E^* \geq y_D$, and (iii) if $m \in (m^*, 1]$, then $y_E^* < y_D$. Moreover, the range $[m_b^*, m^*]$ exists if and only if $\delta \geq \delta^*$, where δ^* is such that $y(m_b^*, \delta^* n, T^*) = y_D$.*

Proposition 4 suggests that the delegation of the decision right to an expert can make the action taken more or less aggressive. We find that taking more or less aggressive action is due to two effects: (1) the *adjustment effect* and (2) the *exclusion effect*. The *adjustment effect* can either make the action more or less aggressive and is due to the decision maker's imperfect knowledge about the state of the random event and the increased efficiency that the expert brings into the decision making process. In particular, the adjustment effect is a result of two potentially conflicting forces: (1a) the increased information processing ability which leads to a more aggressive optimal action and (1b) the elimination of uncertainty which makes the optimal action more or less aggressive.

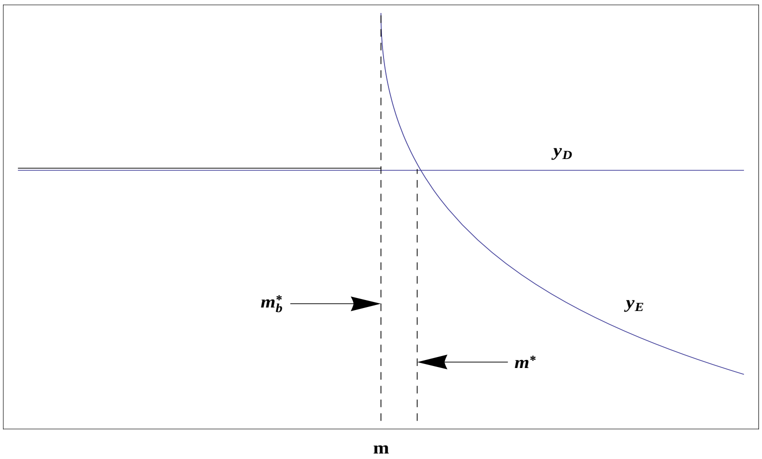

Figure 1: Comparison of the optimal actions taken by the decision maker and the expert as a function of the random event.

When the true state of the random event is low (i.e., m is small), the adjustment effect suggests a more aggressive optimal action since both the processing efficiency and better information lead to an aggressive action. However, when the true state is high, the adjustment effect is ambiguous since higher information processing efficiency leads to more aggressive action but better information leads to a less aggressive one. The adjustment effect suggests a less aggressive action only if the effect of better information dominates the effect of higher information processing. The magnitude of the adjustment effect is measured by $|y_F(m) - y_D|$ when the true state of the random event is m. The *exclusion effect*, on the other hand, always prescribes a less aggressive action and is due to the expert taking advantage of its position and changing the optimal action in a way to increase his utility. The magnitude of the exclusion effect is measured by $|y_F(m) - y_E|$ when the true state of the random event is m.

The following two special cases delineate the drivers of the adjustment and exclusion effects. First, when the expert is used for his information processing ability only (i.e., m is deterministic and $\delta > 1$), the optimal action taken by the expert is always more aggressive than the decision maker's optimal action choice, that is $y_E > y_D$ (see Appendix B for proofs). The increase in the action is entirely due to the increased information processing ability. On the other hand, when the expert is used for elimination of uncertainty only (i.e., m is random and $\delta = 1$), the expert's optimal action recommendation can be higher or lower than the decision maker's optimal action choice. If m is high, then $y_E < y_D$ whereas if m is small, then $y_E > y_D$. In this case, the increase/decrease in the optimal action is entirely due to the effect of better information. Therefore, we can conclude that while the adjustment effect can be driven by either asymmetric information or the expert's ability to superior efficiency in processing noisy information, the exclusion effect is driven by both

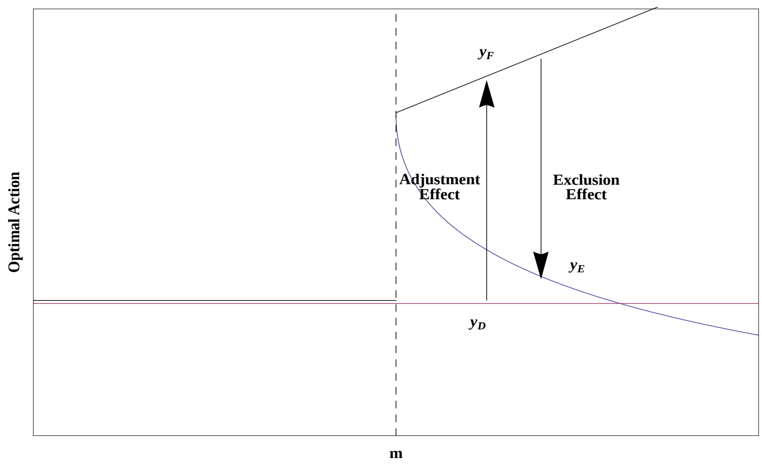

Figure 2: Comparison of the optimal actions taken by the decision maker and the expert as a function of the random event.

effects simultaneously.

Figure 2 illustrates the impact of the adjustment and competitive exclusion effects on the resulting action for $\delta \geq \delta^*$. If the true values of the random event is small, the adjustment effect prescribes a more aggressive action, whereas the exclusion effect demands a less aggressive one. Because the magnitude of the adjustment effect is greater than the magnitude of the exclusion effect ($|y_F(m) - y_D| > |y_F(m) - y_E|$), the net effect is a more aggressive action ($y_E > y_D$). On the other hand, when the random event has a high impact, both the adjustment and the exclusion effects demand a less aggressive action and therefore, the net effect is to take a softer action ($y_E < y_D$). Notice that the adjustment effect reduces the action in this case since, for the chosen parameter set, the action-reduction effect of better information dominates the action-expanding effect of informational efficiency.

To summarize, Proposition 4 suggests that while the expert can lead to reduction in the optimal action, this reduction is not always due to expert's incentives but can also be due to the adjustment effect, which is simply the necessary change for the optimal action to be taken. The adjustment effect can cause a convergence or divergence from the optimal decision, whereas the exclusion effect always causes a divergence.

4 Conclusions

We consider a stylized model where a decision maker relies on an expert for a decision under uncertainty. The goal of our research is to investigate the impact of delegating decision power when it is not possible to write an optimal contract, which in general requires substantial costly analysis. Decision makers benefit from expert's (1) superior knowledge about the nature of the decision and/or (2) efficiency in processing noisy information. Our results are along these two dimensions.

The overall conclusion of our research is that while using experts for decision under uncertainty can be an excellent value proposition for decision makers, the consequences of using experts should be better understood by decision makers. We find that the consequences of using experts may differ depending on what the experts are used for. First, situations where the decision maker needs the expert for costly insights but the uncertainty surrounding the nature of the decision is well understood are perfect candidates for relying on an expert.

Second, in situations where decision makers need insights only, decision makers can use experts who are better in processing noisy information but should be aware that the experts should be rewarded for providing insights. Depending on the value of the random event, relying on an expert may result in an action that is softer or more aggressive when compared to the decision maker's original action decision. The divergence from the decision maker's original action is entirely due to the adjustment effect. Finally, situations where decision makers rely on experts to increase the efficiency in processing noisy information and provide insights related to the decision are also suitable for the use of experts but decision makers need to compensate experts for both his superior information and efficiency. The overall convergence/divergence of the final decision from the optimal decision depends on the magnitudes of the adjustment and exclusion effects.

Appendix

A. Proofs of Propositions and Lemmas

Proof of Proposition 1.

Proof. Because U is linear in m, the decision maker's problem can be rewritten as

$$\max_y E[U(y,m,n)] = U(y,\mu,n)$$

where $\mu = E[m]$. Then, the decision maker maximizes her utility according to first order condition

$$U_1(y,\mu,n) = 0.$$

Because U is concave in y, this first order condition leads to an optimal decision as a function of μ

and n, i.e., $y_D = y(\mu, n)$. By taking the total derivative of the first order condition, we get

$$\begin{aligned} U_{11}dy + U_{12}d\mu + U_{13}dn &= 0 \\ \frac{dy}{d\mu} &= -\frac{U_{12}}{U_{11}} < 0 \\ \frac{dy}{dn} &= -\frac{U_{13}}{U_{11}} > 0 \end{aligned}$$

The inequalities hold since $U_{11} < 0$, $U_{12} < 0$, and $U_{13} > 0$ by assumption. □

Proof of Proposition 2.

Proof. The first order condition is $U_1(y_F, m, \delta n) = 0$, which leads to $y_F = y(m, \delta n)$. By taking the total derivative, we get

$$\begin{aligned} U^D_{11}dy + U^D_{12}dm + nU^D_{13}d\delta &= 0 \\ \frac{dy}{dm} &= -\frac{U^D_{12}}{U^D_{11}} < 0 \\ \frac{dy}{d\delta} &= -n\frac{U^D_{13}}{U^D_{11}} > 0 \end{aligned}$$

The inequalities hold since $U^D_{11} < 0$, $U^D_{12} < 0$, and $U^D_{13} > 0$ by assumption. □

Proof of Lemma 1.

Proof. Let $Y(m, \delta n, T) \equiv \{y : U^D(y, m, \delta n) \geq g(T)\}$ is the set of actions that can generate threshold T as a feasible outcome. That is, $Y(m, \delta n, T)$ determines the set of feasible actions to achieve the threshold T for the expert. Notice that for high enough T, $Y(m, \delta n, T)$ is an empty set since $U^D(y, m, \delta n)$ is concave in y and $g(T)$ is increasing in T. In other words, for each type of the expert, there is a threshold level for which the expert is unable to deliver the target no matter what action he takes.

We first show that $Y(m, \delta n, T)$ is convex and compact. If $Y(m, \delta n, T)$ is an empty set, we are done. So, suppose that $Y(m, \delta n, T)$ is a non-empty set. Consider the special case that $U^D(y, m, \delta) = g(T)$. Because U^D is a concave function of y, there are at most two solutions for this problem. Let $y_1(m, \delta n, T)$ and $y_2(m, \delta n, T)$ be those solutions, i.e., $U^D(y_1, m, \delta n) = U^D(y_2, m, \delta n) = g(T)$. Obviously, $y_1, y_2 \in Y(m, \delta n, T)$. We now show that any $z \in [y_1, y_2]$ also belongs to $Y(m, \delta n, T)$, which is due to the concavity of U^D. Because U^D is concave, we must have

$$U^D(z, m, \delta n) \geq \alpha U^D(y_1, m, \delta n) + (1 - \alpha)U^D(y_2, m, \delta n)$$

where $z = \alpha y_1 + (1-\alpha)y_2$ and $\alpha \in [0,1]$. Because $y_1, y_2 \in Y(m, \delta n, T)$, we have $U^D(y_1, m, \delta n) \geq g(T)$ and $U^D(y_2, m, \delta n) \geq g(T)$, which together with the inequality above imply that $U^D(z, m, \delta n) \geq g(T)$. So, for any $z \in [y_1, y_2]$, it must be true that $z \in Y(m, \delta n, T)$. Notice that this means $Y(m, \delta n, T)$ is a closed range in the real line, which implies that $Y(m, \delta n, T)$ has to be compact. Now, consider the expert's preferred action $\bar{y}_E$, which is assumed to be different than the decision maker's optimal point. If $\bar{y}_E \notin [y_1, y_2]$ then U^E is continuous and monotonic in $[y_1, y_2]$. In that case, Weisstrass Theorem ensures that U^E attains a unique maximum in the compact and convex set $Y(m, \delta n, T)$. In particular, if $\bar{y}_E < y_1$ the expert will choose y_1 as the optimal action and if $\bar{y}_E > y_2$ she will choose y_2. Finally, suppose that $\bar{y}_E \in [y_1, y_2]$. Then, the expert will choose $\bar{y}_E$ rather than y_1 or y_2. Thus, in any case, the expert will choose a unique action.

Now, suppose that the decision maker increases the threshold from T to T'. This will result in a shrink in the range of feasible actions, that is $Y(m, \delta n, T') \subset Y(m, \delta n, T)$ for $T' > T$. If the expert's most preferred action $\bar{y}_E$ is in $Y(m, \delta n, T')$, then $y(m, \delta n, T) = y(m, \delta n, T')$. But if $\bar{y}_E$ is not in $Y(m, \delta n, T')$, then there are two cases: (i) if $\bar{y}_E < y_1(m, \delta n, T')$ then $y(m, \delta n, T) > y(m, \delta n, T')$ and (ii) if $\bar{y}_E > y_2(m, \delta n, T')$ then $y(m, \delta n, T) < y(m, \delta n, T')$. It is straightforward to see that, in all of these cases, the optimal choice of the expert is monotonic in T. $\square$

Proof of Lemma 2.

Proof. By definition of m_b we have $y_E(m_b, \delta n, T) = y_F(m_b, \delta n)$. By taking the total derivative of this equality, we get

$$\frac{\partial y_E}{\partial m_b} dm_b + \frac{\partial y_E}{\partial \delta n}\frac{\partial \delta n}{\partial \delta} d\delta + \frac{\partial y_E}{\partial T} dT = \frac{\partial y_F}{\partial m_b} dm_b + \frac{\partial y_F}{\partial \delta n}\frac{\partial \delta n}{\partial \delta} d\delta$$
$$\left[\frac{\partial y_E}{\partial m_b} - \frac{\partial y_F}{\partial m_b}\right] dm_b + \left[\frac{\partial y_E}{\partial \delta n} - \frac{\partial y_E}{\partial \delta n}\right]\frac{\partial \delta n}{\partial \delta} d\delta + \frac{\partial y_E}{\partial T} dT = 0.$$

Thus, we can write the desired partial derivatives as

$$\frac{dT}{dm_b} = -\left[\frac{\frac{\partial y_E}{\partial m_b} - \frac{\partial y_F}{\partial m_b}}{\frac{\partial y_E}{\partial T}}\right] \quad \text{and} \quad \frac{dT}{d\delta} = -\left[\frac{\frac{\partial y_E}{\partial \delta n} - \frac{\partial y_F}{\partial \delta n}}{\frac{\partial y_E}{\partial T}}\right]\frac{\partial \delta n}{\partial \delta}.$$

Notice that $\frac{\partial y_F}{\partial m_b} = 0$ at m_b by definition. First, suppose that $\bar{y}_E < \bar{y}_D$. Then, by Proposition 2 and Lemma 1, we know that $\frac{\partial y_E}{\partial m_b} \leq 0$, $\frac{\partial y_E}{\partial \delta n} < 0$, $\frac{\partial y_F}{\partial \delta n} > 0$, and $\frac{\partial y_E}{\partial T} > 0$. Thus, $\frac{dT}{dm_b} > 0$ and $\frac{dT}{d\delta} > 0$ when $\bar{y}_E < \bar{y}_D$. Now, suppose that $\bar{y}_E > \bar{y}_D$. By Proposition 2 and Lemma 1, we know that $\frac{\partial y_E}{\partial m_b} \geq 0$, $\frac{\partial y_E}{\partial \delta n} < 0$, $\frac{\partial y_F}{\partial \delta n} > 0$, and $\frac{\partial y_E}{\partial T} > 0$. Thus, $\frac{dT}{dm_b} < 0$ and $\frac{dT}{d\delta} > 0$ when $\bar{y}_E < \bar{y}_D$. $\square$

Proof of Lemma 3.

Proof. Notice first that $T(m_b, \delta n) = U(m_b, \delta n)$ by definition. Thus, the decision maker's optimization problem is equal to

$$\max_{m_b} \quad \int_0^{m_b} U(\mu_b, n) f(m) dm + \int_{m_b}^1 U(m_b, \delta n) f(m) dm$$

By using the Leibniz's Rule, we find that the first order condition of this problem is

$$U(\mu_b, n) f(m_b) + \int_0^{m_b} \frac{\partial U(\mu_b, n)}{\partial \mu_b} \frac{\partial \mu_b}{\partial m_b} f(m) dm - U(m_b, \delta n) f(m_b) + \frac{\partial U(m_b, \delta n)}{\partial m_b} [1 - F(m_b)] = 0$$

and it determines the borderline type x_b^* that is optimal for the retailer. Notice that $\mu_b = E[m|m \leq m_b] = \int_0^{m_b} m f(m) dm$ by definition and $\frac{\partial \mu_b}{\partial m_b} = m_b f(m_b)$. By rearranging the first order condition, we get

$$\frac{\partial U(m_b, \delta n)}{\partial m_b} = \lambda(m_b) \left[U(m_b, \delta n) - U(\mu_b, n) - \frac{\partial U(\mu_b, n)}{\partial \mu_b} m_b f(m_b) \right]$$

where $\lambda(m_b) = \frac{f(m_b)}{1 - F(m_b)}$. $\square$

Proof of Proposition 3.

Proof. Suppose that the decision maker's beliefs are specified by $\xi^*(m)$. By Lemma 3, we know that it is sequentially rational for the decision maker to offer the target $T^* = T(m_b^*, \delta n)$. Notice that the expert has always incentives to agree with the decision maker since his outside option is zero. Notice also that the expert whose type is m will never accept the offer if $m < m_b^*$ since $T(m, \delta n) < T(m_b^*, \delta n)$ by Lemma 2 and not delivering the target is prohibitively costly. Thus, it is sequentially rational for the expert to agree with the offer when $m \geq m_b^*$ and reject it when $m < m_b^*$. Now, suppose that the players follow their sequentially rational strategies. Then, whenever the decision maker see a rejection, she will infer that the expert cannot be a type $m > m_b^*$. Thus, the decision maker updates her beliefs according to the Bayes' rule, which gives us the belief structure specified by $\xi^*(m)$. $\square$

Proof of Proposition 4.

Proof. First, notice that if $m < m_b^*$, then the expert rejects the offer and the decision maker updates her beliefs and takes the action as in the case where she decides without the help of the expert; which implies that $y_E^* = y(\mu_b, n) > y(\mu, n) = y_D$.

Now, suppose that $m \geq m_b^*$. Notice that the decision maker will never choose a target that is smaller than her utility when she takes the action without the help of the expert; that is, $T(m_b^*, \delta n) \geq U(\mu, n)$. Then, $y_E^* < y_D$ when the expert's type is m_b^* and $\delta = 1$. This implies that

$y_E^* < y_D$ for all $m \geq m_b^*$. Let m^* be a type such that $y(m^*, \delta n) = y_D$. Notice that m^* does not have to be in the range $[0, 1]$. We know that $y(\cdot)$ is increasing in δ. Because $y(m^*, \delta n) = y_D$ and $y(\cdot)$ is decreasing in m^*, we must have $\frac{\partial m^*}{\partial \delta} > 0$. Because $y_E^* < y_D$ when $\delta = 1$, there exists a δ^* such that $y_E^* = y_D$ for m_b^*, the minimum accepting type. Thus, the range $[m_b^*, m^*]$ exists only if $\delta \geq \delta^*$. Also, if $m \in [m_b^*, m^*]$, then $y_E^* \geq y_D = y(m^*, \delta n)$ for any $\delta \geq \delta^*$. Similarly, if $m \in (m^*, 1]$, then $y_E^* < y_D = y(m^*, \delta n)$ since $y(\cdot)$ is decreasing in m. $\square$