NBER WORKING PAPER SERIES

WHEN CONSENSUS CHOICE DOMINATES INDIVIDUALISM:
JENSEN'S INEQUALITY AND COLLECTIVE DECISIONS UNDER UNCERTAINTY

Charles F. Manski

Working Paper 15172
http://www.nber.org/papers/w15172

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2009

This research was supported in part by National Science Foundation grant SES-0911181. I am grateful for comments from Stephen Coate, Hanming Fang, and Aleksey Tetenov. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

WHEN CONSENSUS CHOICE DOMINATES INDIVIDUALISM: Jensen's Inequality and Collective Decisions under Uncertainty
Charles F. Manski
NBER Working Paper No. 15172
July 2009
JEL No. D7,D81,H42

ABSTRACT

Research on collective provision of private goods has focused on distributional considerations. This paper studies a class of problems of decision under uncertainty in which the argument for collective choice emerges from the mathematics of aggregating individual payoffs. Consider decision making when each member of a population has the same objective function, which depends on an unknown state of nature. If agents knew the state of nature, they would make the same decision. However, they may have different beliefs or may use different decision criteria. Hence, they may choose different actions even though they share the same objective. Let the set of feasible actions be convex and the objective function be concave in actions, for all states of nature. Then Jensen's inequality implies that consensus choice of the mean privately-chosen action yields a larger aggregate payoff than does individualistic decision making, in all states of nature. If payoffs are transferable, the aggregate payoff from consensus choice may be allocated to Pareto dominate individualistic decision making, in all states of nature. I develop these ideas. I also use Jensen's inequality to show that a planner with the power to assign actions to the members of the population should not diversify. Finally, I give a version of the collective choice result that holds with consensus choice of the median rather than mean action.

Charles F. Manski
Department of Economics
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
cfmanski@northwestern.edu

1. Introduction

Economists working in the utilitarian paradigm have a strong predisposition to think that individualistic decision making regarding private goods is superior to collective choice. Research on public provision of private goods, whether normative or positive, has largely focused on distributional considerations. See, for example, Blackorby and Donaldson (1988), Besley and Coate (1991), Bruce and Waldman (1991), Blomquist and Christiansen (1995), Coate (1995), and Epple and Romano (1996). A recent exception is Fang and Norman (2008), where government observation of demand for publicly provided goods reveals private information useful in setting tax policy.

This paper gives an entirely different utilitarian rationale for collective choice of private goods. I study a simple problem of private decision making under uncertainty in which the argument for collective choice emerges directly from the mathematics of aggregating individual payoffs.

I mainly consider decision making when each member of a population faces the same choice problem and wants to optimize the same objective function, which depends on an unknown state of nature. If agents knew the state of nature, they would make the same decision. However, they may have different beliefs about the state or may cope with their incomplete information by using different criteria to make decisions. Hence, acting independently, they may choose different actions even though they share the same objective.

For example, a set of price-taking firms may have the same production technology and want to choose inputs to maximize profit. When inputs are chosen, firms may not know the price they will receive per unit of output. Then firms with different beliefs about product price or different risk preferences may choose different input levels.

Suppose that the set of feasible actions is convex and the objective function is concave in actions, for all states of nature. In this context, Section 2 shows that population coordination on a specific consensus action yields a larger aggregate payoff than does individualistic decision making, in all states of nature. The consensus action with this remarkable property is the mean action that agents would choose independently. This result is an immediate consequence of Jensen's inequality. I use input choice by firms and asset allocation by investors to illustrate the result. I discuss its robustness to relaxation of the assumption that each member of the population wants to optimize the same concave objective function.

Section 3 shows that, if payoffs are transferable, the aggregate payoff realized by collective choice of the consensus action may be allocated across the population so that collective choice Pareto dominates individualistic decision making, in all states of nature. A Pareto dominating collective choice mechanism is implementable if agents truthfully reveal the actions they would choose individualistically. I give conditions under which truthful revelation is incentive compatible. I suggest that the collective choice mechanism introduced here may find profitable application in the operation of agricultural cooperatives.

Whereas Section 3 concerns decision making by a population whose members must agree on a collective choice mechanism, Section 4 considers treatment choice by a planner who has the power to assign actions to the members of the population. Here Jensen's inequality implies that any heterogeneous treatment of the population is dominated by assigning the associated mean treatment to all persons. Thus, the planner should not diversify treatment. I explain why this result differs from the positive findings for diversification obtained in Manski (2009). There I studied a different class of planning problems, where the choice set is not convex and agents may have heterogeneous, non-concave objective functions.

Except for the discussion of robustness in Section 2, the paper thus far concerns decision making with a convex choice set and a common concave objective function. Section 5 supposes instead that the choice set is ordered and that agents have objective functions that are unimodal in each state of nature. The state-specific mode is invariant across agents, but the objective functions may otherwise vary. In this setting, I show that consensus choice of the median individualistic action makes the majority of agents better off than they would be with private choice, in all states of nature. This result is closely related to the median voter theorem.

The analysis of this paper is simple and straightforward. It nevertheless appears to be new to research on collective choice of private goods. The closest connection that I have been able to discover is to the distantly related literature on consensus forecasting. Section 6 makes this connection.

2. Analysis with a Convex Choice Set and a Concave Objective Function

2.1. Notation and Concepts

To begin, let (J, Ω, P) be a probability space of agents, each of whom must choose an action from a set X. Here J lists the agents, the σ-algebra Ω places probability on individual agents, and P is the probability measure. For example, J may be a finite group of size $|J|$, in which case $P(j) = 1/|J|$ for $j \in J$. Or J may be a continuum indexed by the unit interval, in which case P is the uniform distribution on [0, 1].

The set S lists feasible states of nature. The objective function $f(\cdot, \cdot): X \times S \rightarrow R$ maps actions and states into the real line. Each agent wants to choose an action that maximizes $f(\cdot, r)$, where $r \in S$ is the actual state of nature. Agents know that $r \in S$, but they may not know the identity of r. Hence, they may not be able to solve the optimization problem $\max_{x \in X} f(x, r)$.

Agents may have different beliefs about r. They may also use different criteria to make choices with their beliefs. For example, they may maximize subjective expected utility or apply the maximin or minimax-regret criterion. Let x_j be the action that agent $j \in J$ would choose in a regime of individualistic decision making. Then the aggregate payoff in state of nature s is $\int f(x_j, s)dP(j)$.

Consider an alternative regime of collective choice in which all agents choose some consensus action, say $c \in X$. Then the aggregate payoff in state s is $\int f(c, s)dP(j) = f(c, s)$. The question of interest is whether there exists a consensus choice that improves on individualistic decision making. The answer is positive if X is convex and if $f(\cdot, s)$ is concave on X for each $s \in S$.

2.2. Collective Choice of the Mean Action

Let $\mu \equiv \int x_j dP(j)$ denote the mean individualistic action and assume that μ is finite. Convexity of X implies that $\mu \in X$, so μ is a feasible action. The aggregate payoff with collective choice of μ is $f(\mu, s)$. Jensen's inequality gives

$$(1) \quad f(\mu, s) \geq \int f(x_j, s)dP(j), \quad \text{all } s \in S.$$

Thus, in all states of nature, the aggregate payoff with consensus choice of action μ is at least as large as with individualistic decision making.

Let $g(s) \equiv f(\mu, s) - \int f(x_j, s)dP(j)$ denote the aggregate surplus achieved by consensus choice of μ. This surplus is non-negative by (1). It ordinarily is positive, a familiar sufficient condition being that $f(\cdot, s)$ is strictly concave and $Var(x_j) > 0$.

2.3. Illustrative Applications to Production Decisions

Result (1) has many applications in which a group of firms or other agents have a common concave production function and want to choose inputs to maximize profit. Here are two examples.

Production by Firms when Product Price is Unknown

Let J be a group of price-taking firms which face the same concave production function. In particular, let the set of feasible inputs be $X = [0, \infty)$ and let output be $\log(x + 1)$. Let the unit cost of input be one, and let product price be s. Then the profit function is $f(x, s) = s\cdot\log(x + 1) - x$. Suppose that firms do not know product price when they choose input quantities. For example, the firms may be farms of equal size which must decide how intensively to plant in the spring, before knowing the crop price that they will receive at harvest.

Let the set S of feasible prices be a subset of the half-line $[0, \infty)$. Suppose each firm places a subjective distribution on price and maximizes expected profit. Let p_j be the subjective mean price held by firm j. Acting independently, firm j would solve the problem $\max_{x \in X} p_j\cdot\log(x + 1) - x$, yielding the input choice $x_j = \max(0, p_j - 1)$. Suppose that $p_j \geq 1$ for all $j \in J$. Then $\mu = E(p) - 1$.

In state s, the aggregate payoff with collective choice of input quantity μ and with individualistic decision making are respectively

$$f(\mu, s) \;=\; s\cdot\log[E(p)] - E(p) + 1, \qquad\qquad \int f(x_j, s)dP(j) \;=\; sE[\log(p)] - E(p) + 1.$$

Hence, the surplus achieved by consensus choice of μ is $g(s) = s\{\log[E(p)] - E[\log(p)]\}$.

Allocation of an Endowment Between a Safe Asset and Production with Unknown Return

Let J be a group of investors, each having an endowment of size one. Suppose that each investor must allocate his endowment between a safe asset with known return one and a productive activity with unknown return. The set of feasible choices is $X = [0, 1]$, where $x \in X$ denotes the fraction of the endowment allocated to the productive activity. Let the return to production be $s\cdot x^{1/2}$, where $S = (0, \infty)$. Then the return to allocation x in state of nature s is $s\cdot x^{1/2} + (1 - x)$.

Suppose each investor places a subjective distribution on S and maximizes expected return. Let q_j be the subjective mean of s for investor j. Acting independently, j would solve the problem $\max_{x \in X} q_j\cdot x^{1/2} + (1 - x)$, yielding the choice $x_j = \min(1, q_j^2/4)$. Suppose that $q_j \leq 2$ for all $j \in J$. Then $\mu = E(q^2)/4$.

In state s, the aggregate payoff with collective choice of input quantity μ and with individualistic decision making are respectively

$$f(\mu, s) \;= (s/2)\cdot[E(q^2)]^{1/2} + 1 - E(q^2)/4, \qquad\qquad \int f(x_j, s)dP(j) \;= (s/2)\cdot E(q) + 1 - E(q^2)/4.$$

Hence, the surplus achieved by consensus choice of μ is $g(s) = (s/2)[E(q^2)^{1/2} - E(q)]$.

2.4. Relaxing the Assumption of a Common Concave Objective Function

Observe that result (1) requires no restrictions on the beliefs that agents hold about the actual state of nature, nor on the decision criteria they use to cope with incomplete knowledge of the state. Some agents may believe that they know the state of nature with certainty. Others may have probabilistic beliefs and maximize expected utility, with possibly heterogeneous risk preferences. Still others may not have probabilistic beliefs and use a criterion for decision making under ambiguity, such as the maximin or minimax-regret criterion.

The result as stated does presume that each member of the population has the same concave objective function. Although some semblance of this condition seems essential, some relaxation is possible. I consider two scenarios here.

Collective Choice within Sub-populations

Let the population be composed of K types. Persons of a given type have the same concave objective function, but this function may vary across types. Let μ_k denote the mean action that would be privately chosen by persons of type k. Consider collective choice in sub-populations, with all persons of type k choosing action μ_k.

Application of (1) to persons of type k gives $f(\mu_k, s) \geq \int f(x_j, s)dP(j|k)$, all $s \in S$. Aggregation of payoffs across types gives $\sum_k f(\mu_k, s)P(k) \geq \int f(x_j, s)dP(j)$, where P(k) is the fraction of type-k persons in the population. Thus, the aggregate payoff from collective choice within each of the K sub-populations is at least as large as that produced by individualistic decision making.

The above proves that collective choice within each of the K sub-populations improves on any decision process based on a finer partition of the population. It does not prove that the aggregate payoff is maximized by partitioning the population into the K types. It leaves open the possibility that collective choice within a grosser partition of the population would further increase the aggregate payoff in at least some states of nature.

Neighborhoods of a Common Concave Function

Suppose that each agent $j \in J$ has a person-specific, not necessarily concave, objective function $f_j(\cdot, \cdot)$: $X \times S \rightarrow R$. Then the aggregate payoffs in state of nature s with individualistic decision making and collective choice of μ are $\int f_j(x_j, s)dP(j)$ and $\int f_j(\mu, s)dP(j)$ respectively. Collective choice achieves a non-negative surplus if $\int f_j(\mu, s)dP(j) \geq f_j(x_j, s)dP(j)$. This inequality can hold if agents do not share the same concave objective function. In particular, it holds if all agents have objective functions that are sufficiently close to a common concave function.

For each state s, let there exist a $\lambda(s) > 0$ and a concave function $f(\cdot, s)$: $X \rightarrow R$ such that

$$\sup_{w \in X, j \in J} |f_j(w, s) - f(w, s)| < \lambda(s).$$

As earlier, let $g(s) \equiv f(\mu, s) - \int f(x_j, s)dP(j)$ denote the surplus that occurs with function f. Surplus with the actual objective functions $f_j(\cdot, s)$, $j \in J$ satisfies the inequality

$$\int f_j(\mu, s)dP(j) - \int f_j(x_j, s)dP(j) \geq g(s) - 2\lambda(s).$$

Hence, the condition $\lambda(s) \leq g(s)/2$ suffices to ensure that actual surplus is non-negative.

3. Pareto Dominant Collective Choice Mechanisms

3.1. Allocation of Transferable Payoffs

The fact that collective choice of μ yields non-negative aggregate surplus in all states of nature does not imply that collective choice makes all agents better off relative to individualistic decision making. However, if payoffs are transferable across the population, then there exist collective choice mechanisms that Pareto dominate private choice.

Let $\gamma \equiv [\gamma_j(s), (j, s) \in J \times S]$ be any set of positive real numbers such that $\int \gamma_j(s)dP(j) \leq 1$, all $s \in S$. Consider a collective choice mechanism in which agents take the consensus action μ and, if the state of nature turns out to be s, agent j receives the payoff

(2) $h_j(\gamma, s) \equiv f(x_j, s) + \gamma_j(s)g(s).$

The collective-choice payoffs weakly dominate payoffs with private choice; that is, $h_j(\gamma, s) \geq f(x_j, s)$ for all $s \in S$ and $h_j(\gamma, s) \geq f(x_j, s)$ for some s. The collective-choice payoffs are feasible because $\int h_j(\gamma, s)dP(j) \leq f(\mu, s)$. Hence, the collective choice mechanism is feasible and Pareto dominates individualistic decision making. Selecting a γ with $\int \gamma_j(s)dP(j) = 1$ allocates the surplus fully in every state of nature.

Observe that a collective choice mechanism of this type benefits even an agent who knows the actual state of nature with certainty. The consensus action μ may differ from what such an agent knows to be the optimal action. Nonetheless, in all states of nature, he receives at least the payoff that he would have achieved by private choice.

Our finding that consensus choice of a private good can Pareto dominate individualistic decision making goes against conventional economic wisdom. It has been common to think that, absent distributional considerations or paternalism, private choice must be Pareto superior to collective choice. The standard utilitarian argument is that when agents want to make different choices, collective choice of a single good creates deadweight loss.

Our analysis shows that it is important to ask why agents want to make different choices. The standard utilitarian argument is correct in deterministic settings, where variation in private choices stems solely from heterogeneity in choice sets or objective functions. In decisions under uncertainty, variation in private choices may also stem from heterogeneity in agents' beliefs about the state of nature or in their decision criteria. The latter sources of heterogeneity are the driving force behind result (1). When agents have the same convex choice set and concave objective function, (1) shows that consensus choice creates aggregate surplus rather than deadweight loss. This enables the Pareto improvement achieved by mechanism (2).

Contrast with Collective Insurance

The collective choice mechanism introduced here differs from collective insurance. In insurance systems, agents make individualistic decisions and then smooth payoffs across states of nature. Insurance increases payoffs in states with relatively bad private outcomes and decreases payoffs in states with relatively good private outcomes. Typically, agent j privately chooses an action x_j and receives payoff $f(x_j, s) + \varphi_j(s) - c_j$. Here $c_j > 0$ is a state-invariant insurance premium and $\varphi_j(s)$ is a state-dependent insurance payment, with $\varphi_j(s) = 0$ in states where $f(x_j, s)$ is relatively large and $\varphi_j(s) > c_j$ in states where $f(x_j, s)$ is relatively small.

A collective insurance system is feasible if aggregate payments do not exceed aggregate premiums in every state of nature; that is, if $\int \varphi_j(s)dP(j) \leq \int c_j dP(j)$ for all $s \in S$. This set of inequalities formalizes the idea that it is infeasible to insure against systemic risks. Insurance is feasible only if, in every state of nature, the sub-population of agents who experience bad private outcomes and receive payments is sufficiently small that aggregate premiums cover their payments.

In contrast to insurance, our agents all choose the consensus action μ. In every state of nature, they receive payoffs (2) that are at least as large as their private outcomes under individualistic decision making. These payoffs are feasible because collective choice of μ generates a non-negative aggregate surplus in every state of nature.

3.2. Truthful Revelation of Private Choices

To implement the collective choice mechanism requires knowledge of the actions that agents would choose privately. Research on mechanism design has long sought to determine when it is incentive compatible for agents to reveal private information truthfully. Here we would like each agent to announce his private choice truthfully.

In a regime of private choice, agent j receives payoff $f(x_j, s)$ in state s. Let x_j^a denote the action that j announces under the collective choice mechanism and let $\mu^a \equiv \int x_j^a dP(j)$ denote the mean of the announced actions. Then

$$(2') \quad h^a_j(\gamma, s) \equiv f(x^a_j, s) + \gamma_j(s)g^a(s)$$

is the collective-choice payoff based on announcements. Here $g^a(s) \equiv f(\mu^a, s) - \int f(x^a_j, s)dP(j)$ is the surplus in state s.

Private and collective-choice payoffs differ, so agent j would not necessarily announce $x^a_j = x_j$. However, truthful revelation is incentive compatible in some settings. I show here that it is if the population is a continuum and agents either maximize expected payoff or minimize maximum regret.

Suppose that the population is a continuum. Then the action announced by agent j does not affect $g^a(s)$. Here, as in analysis of other collective decision problems, the idea of a continuum of agents is a simplifying idealization, meant to approximate a large finite population. In a large finite population, an agent's announced action negligibly affects aggregate surplus.

Suppose that j announces an action that maximizes his subjective expected payoff. Then $x^a_j = \operatorname{argmax}_{x \in X} \int f(x, s)d\pi_j + \int \gamma_j(s)g^a(s)d\pi_j$, where π_j is the agent's subjective distribution on S. In the private-choice regime, this agent would choose $x_j = \operatorname{argmax}_{x \in X} \int f(x, s)d\pi_j$. The expression $\int \gamma_j(s)g^a(s)d\pi_j$ does not vary with x. Hence, $x^a_j = x_j$.

Suppose that j minimizes maximum regret. In this case $x^a_j = x_j$ because

$$\max_{w \in X} [f(w, s) + \gamma_j(s)g^a(s)] - [f(x, s) + \gamma_j(s)g^a(s)] = \max_{w \in X} f(w, s) - f(x, s).$$

The left-hand side of the equation is regret for announcing action x in state s. The right-hand side is regret for choosing x in state s in a regime of private choose. The equation holds because, the population being a continuum, $g^a(s)$ does not vary with w or x.

3.3. Application to Agricultural Cooperatives

Is the collective choice mechanism introduced here only of theoretical interest, or might it have useful application? The mechanism may prove particularly helpful to stimulate the formation and guide the operation of agricultural cooperatives. I use crop production to illustrate.

Crop production approximates the conditions assumed in this paper reasonably closely. Farmers share the common objective of profit maximization. They typically have concave production functions and convex choice sets, with diminishing returns to inputs of seed, fertilizer, labor, and irrigation. The farmers in a given region face approximately homogeneous climate, soil conditions, input costs, and crop prices. Hence, their profit functions are approximately homogeneous. Finally, profits are a transferable payoff. Thus, the conditions for application of result (1) and mechanism (2) approximately hold.

The surplus generated by consensus choice is greatest in settings with considerable uncertainty, where agents with disperse beliefs and decision criteria make heterogeneous private choices. Farmers making planting decisions in the spring face considerable uncertainty about the crop prices they will receive in the fall, as well as uncertainty about crop yield arising from the difficulty of predicting weather during the planting season. In a private-choice regime, farmers may cope with this uncertainty in different ways, varying the allocation of fields to alternative crops, the timing of planting, and so on. Hence, consensus choice may generate meaningful surplus.

A cooperative could determine the consensus input bundle in late winter, prior to initiation of the annual production cycle in the spring. Each farmer in the cooperative would report the input bundle he would choose if making production decisions individualistically. The mean announced bundle would become the consensus choice.

4. Implications for Planning: Uniform Treatment Dominates Diversification

Sections 2 and 3 considered decision making when the members of a population collectively agree on a consensus choice. This section concerns a planner who has the power to assign actions. The planner's objective is to maximize the aggregate payoff.

When studying planning problems, actions are often called treatments. The planner's problem is to choose treatments. A planner with incomplete knowledge of the state of nature has partial knowledge of treatment response. I have previously studied certain such planning problems in Manski (2005, 2009) and elsewhere. Those problems did not have a convex choice set and concave objective function. I show here that this structure qualitatively affects findings.

Suppose that a planner can assign each agent any feasible action. Thus, the planner can choose any element of the Cartesian Product set $X^{|J|}$. Let $w \equiv (w_j, j \in J)$ be any set of assigned actions with finite mean $\mu_w \equiv \int w_j dP(j)$. Jensen's inequality gives

$$(3) \quad f(\mu_w, s) \;\geq\; \int f(w_j, s)dP(j), \quad \text{all } s \in S.$$

Result (3) shows that, in each state, the aggregate payoff when the planner assigns every agent action

μ_w is at least as large as the payoff with the possibly heterogeneous assignments (w_j, $j \in J$). In other words, any diversified treatment of the population is dominated by assigning the associated mean treatment to all persons.

This finding differs sharply from one that I have reported in Manski (2009) and elsewhere. There I studied a class of planning problems in which diversified treatment of observationally identical persons is always undominated. In particular, I studied problems where the choice set X contains only two elements, say a and b. I also permitted heterogeneous objective functions.

Let $\alpha(s) \equiv \int f_j(a, s)dP(j)$ denote the aggregate payoff in state s when all agents receive treatment a and, similarly, let $\beta(s) \equiv \int f_j(b, s)dP(j)$. Consider a treatment allocation that randomly assigns a fraction $\delta \in [0, 1]$ of the population to treatment b and $1 - \delta$ to treatment a. Let the population be a continuum. Then the payoff to allocation δ in state s is $(1 - \delta)\alpha(s) + \delta\beta(s)$. Suppose that the planner faces ambiguity; that is, there exists a state s with $\alpha(s) > \beta(s)$ and another state s' with $\alpha(s') < \beta(s')$. Then it is immediate that all $\delta \in [0, 1]$ are undominated. Moreover, it can be shown that the minimax-regret allocation always diversifies treatment.

Result (3) shows that the situation differs dramatically if all convex combinations of a and b are feasible treatments and all agents have the same concave objective function f. Then, for all $\delta \in (0, 1)$ and $s \in S$, assigning every agent to treatment $(1 - \delta)a + \delta b$ yields at least as large a payoff as does the diversified allocation assigning fraction δ to b and $1 - \delta$ to a. Formally,

$$f[(1 - \delta)a + \delta b, s] \;\geq\; (1 - \delta)f(a, s) + \delta f(b, s) \;=\; (1 - \delta)\alpha(s) + \delta\beta(s), \text{ all } s \in S.$$

Application to Medical Treatment

Medical treatment with partial knowledge of treatment response illustrates when diversification is and is not a reasonable strategy. Consider first an organ disease with two alternative treatments. One is surgery to repair the organ and the other is replacement of the organ with a transplant. Convex combinations of these treatments are not feasible—one can only repair or replace. In a setting of this sort, diversification warrants consideration when it is not clear which treatment is better. Some fraction of patients would have the organ repaired and the remaining fraction would receive transplants. The minimax-regret criterion provides a coherent method to choose the fractions.

Now consider a psychiatric ailment where the potential treatments are alternative doses of a certain drug and the outcome of interest is a real measure of patient functionality. Here convex combinations of treatments are feasible—one can administer the drug in any dose. Suppose the medical planner knows that the objective function is concave and homogeneous across the patient population, with diminishing marginal returns to larger doses and possibly decreased functionality beyond some dosage. In a setting of this sort, a medical planner with partial knowledge of the dose response function should not vary dosage across patients. Any diversified treatment strategy is dominated by a decision to give all patients the mean of the contemplated doses.

5. Analysis with An Ordered Choice Set and Unimodal Objective Functions

Section 2.4 showed that collective choice of the mean action may dominate individualistic decision making in circumstances where agents do not have the same concave objective function. However, this goes so only so far. Result (1) certainly does not hold universally.

This section presents a weaker result that holds when X is an ordered set of actions and each agent j has an objective function $f_j(\cdot, \cdot)$: $X \times S \rightarrow R$ that is unimodal on X for each $s \in S$. Set X need not be convex; it may, for example, be a discrete subset of the real line. Objective functions may vary across agents, except that they share the same mode. Formally, I assume that, for each $s \in S$, $\theta(s) \equiv \text{argmax}_{x \in X}\, f_j(x, s)$ is invariant across $j \in J$.

Let m denote the median individualistic action; that is, $m \equiv \inf t: P(x_j \leq t) \geq ½$. I will show that, in every state of nature, a majority of agents receive at least as high a payoff with consensus choice of action m as they would with individualistic decision making. That is,

(4) $P[f_j(m, s) \geq f_j(x_j, s)] \geq ½$, all $s \in S$.

Proof: Result (4) is a special case of a result that holds for any quantile of the distribution of individualistic actions. For $\kappa \in (0, 1)$, let q_κ denote the κ-quantile individualistic action; that is, $q_\kappa \equiv \inf t: P(x_j \leq t) \geq \kappa$. The general result is

(5) $P[f_j(q_\kappa, s) \geq f_j(x_j, s)] \geq \min(\kappa, 1 - \kappa)$, all $s \in S$.

To show (5), consider separately the cases in which $q_\kappa = \theta(s)$, $q_\kappa < \theta(s)$, and $q_\kappa > \theta(s)$.

First let $q_\kappa = \theta(s)$. Then $P[f_j(q_\kappa, s) \geq f_j(x_j, s)] = 1$. Next let $q_\kappa < \theta(s)$. Then $P[f_j(q_\kappa, s) \geq f_j(x_j, s)] \geq P(x_j \leq q_\kappa) \geq \kappa$. Now let $q_\kappa > \theta(s)$. Then $P[f_j(q_\kappa , s) \geq f_j(x_j, s)] \geq P(x_j \geq q_\kappa) \geq 1 - \kappa$. Combining these cases yields (5). Result (4) is the special case with $\kappa = \frac{1}{2}$. □

Result (4) brings to mind the median voter theorem of Black (1948), which also considered an ordered set of actions and unimodal objective functions. Indeed (4) shows that, given knowledge of the state of nature, a majority of agents prefer consensus choice of action m to the set $(x_j, j \in J)$ of individualistic actions. However, (4) does not imply that a majority of agents prefer m ex ante, before the actual state is known. The composition of the majority who receive larger payoffs with consensus choice varies with the state of nature. Hence, agents need not prefer m ex ante.

While result (4) is interesting, it is much less powerful than result (1). It does not provide the foundation for a Pareto dominant collective choice mechanism. The aggregate surplus in state s with collective choice of m is $\int f_j(m, s)dP(j) - \int f_j(x_j, s)dP(j)$. Result (4) does not determine the sign of this quantity. If there exists a state of nature where surplus is negative, it is not possible in this state to allocate payoffs so all agents are better off than they would be with private choice.

6. Jensen's Inequality and Research on Consensus Forecasting

The analysis of this paper appears to be entirely new to research on collective choice of private goods. However, a version of result (1) has received sporadic recognition in research on the distantly related subject of consensus forecasts. I explain in this concluding section.

For over a century, beginning at least as early as Galton (1907), researchers studying the accuracy of forecasts have studied settings in which multiple agents are asked to give point forecasts of a quantity and their forecasts are combined to create a *consensus forecast*. It is particularly common to define the consensus forecast as the cross-sectional mean of the individual forecasts.

Empirical studies have regularly found that the cross-sectional mean forecast is more accurate than the individuals forecasts used to form the mean. Clemen (1989, p. 559) put it this way in a review article:

> "The results have been virtually unanimous: combining multiple forecasts leads to increased forecast accuracy. This has been the result whether the forecasts are judgmental or statistical, econometric or extrapolation. Furthermore, in many cases one can make dramatic performance improvements by simply averaging the forecasts."

Researchers have been intrigued by this empirical regularity, which has recently become known popularly as the "wisdom of crowds" (Surowiecki, 2004). Various reasons have been suggested. In fact, the empirical regularity follow from Jensen's inequality.

Let y_n, n = 1, . . . , N be a set of individual point forecasts of an unknown real quantity θ, let P_N denote the cross-sectional distribution of the forecasts, and let $\mu_N \equiv \int ydP_N$ denote the cross-sectional mean forecast. Let $L(\cdot, \cdot): R \times R \rightarrow [0, \infty)$ be a loss function used to measure the consequence of prediction error. Research on forecasting has typically used absolute loss $L(y, \theta) = |y - \theta|$ or square loss $L(y, \theta) = (y - \theta)^2$. When these or any other convex loss function is used, Jensen's inequality gives $L(\mu_N, \theta) \leq \int L(y, \theta)dP_N$ for all $\theta \in R$. Thus, whatever the actual value of the quantity being forecast, the loss associated with the mean forecast is no larger than the mean loss of the individual forecasts.

This simple result has long been known in statistical decision theory. There θ is a parameter to be estimated and (y_n, n = 1, . . . , N) is a *randomized estimate*, meaning that the statistician draws an integer i at random from the set (1, . . . , N) and uses y_i to estimate θ. Suppose that a convex loss function is used to measure precision of estimation. Then Jensen's inequality implies that loss using the *non-randomized estimate* μ_N is smaller than expected loss using the randomized estimate. See Hodges and Lehmann (1950).

Research on consensus forecasts has largely disregarded the result as it has sought to explain why mean forecasts perform better than individual forecasts. A notable exception is McNees (1992), who exposited the matter clearly in the context of absolute and square loss. He also recognized the consensus-forecasting version of result (4); that is, the median forecast of any event must be at least as close to the truth as at least half of the individual forecasts, whatever the truth may be. McNees observed that much research on forecasting did not acknowledge "these simple, well-known, yet often ignored arithmetic principles" (page 705).

More recently, Larrick and Soll (2006) referred to the application of Jensen's inequality to consensus forecasting as the "averaging principle" and reported experimental research showing that a majority of their student subjects did not understand the principle. Thus, the power of Jensen's inequality may be plain in abstraction but not as evident in application.

References

Besley, T. and S. Coate (1991), "Public Provision of Private Goods and the Redistribution of Income," *American Economic Review*, 81, 979-984.

Black, D. (1948), "On the Rationale of Group Decision-Making," *Journal of Political Economy*, 56, 23-34.

Blackorby, C. and D. Donaldson (1988), "Cash versus Kind, Self-Selection, and Efficient Transfers," *American Economic Review*, 78, 691-700.

Blomquist, S. and V. Christiansen (1995), "Public Provision of Private Goods as a Redistributive Device in an Optimum Income Tax Model," *Scandinavian Journal of Economics*, 97, 547-67.

Bruce, N. and M. Waldman (1991), "Transfers In-Kind: Why They Can Be Efficient and Nonpaternalistic," *American Economic Review*, 81, 1345-1351.

Coate, S. (1995), Altruism, the Samaritan's Dilemma, and Government Transfer Policy," *American Economic Review*, 85, 46-57.

Clemen, R. (1989), "Combining Forecasts: A Review and Annotated Bibliography," *International Journal of Forecasting,* 5, 559-583.

Epple, D. and R. Romano (1996), "Public Provision of Private Goods," *Journal of Political Economy*, 104, 57-84.

Fang, H. and P. Norman (2008), "Toward an Efficiency Rationale for the Public Provision of Private Goods," National Bureau of Economic Research Working Paper 13827.

Galton, F. (1907), "Vox Populi," *Nature*, 75, 450-451.

Hodges, J. and E. Lehmann (1950), "Some Problems in Minimax Point Estimation," *Annals of Mathematical Statistics*, 21, 182-197.

Larrick, R. and J. Soll (2006), "Intuitions about Combining Opinions: Misappreciation of the Averaging Principle," *Management Science*, 52, 111-127.

Manski, C. (2005), *Social Choice with Partial Knowledge of Treatment Response*, Princeton: Princeton University Press.

Manski, C. (2009), "Diversified Treatment under Ambiguity," *International Economic Review*, forthcoming.

McNees, S. (1992), "The Uses and Abuses of 'Consensus' Forecasts," *Journal of Forecasting*, 11, 703-710.

Surowiecki, J. (2004), *The Wisdom of Crowds*, New York: Random House.