Reciprocity, Proportionality, and Incentives

Ehud Guttel and Shmuel Leshem

September 6, 2013

Abstract

When multiple agents engage in a joint project, allocation of the output—either product or harm—should respect fairness and provide optimal incentives. This paper presents such an effort-based sharing rule for allocating output between identical agents. Under a Reciprocity-Based (RB) sharing rule, (i) one agent is assigned one half of the output that would have resulted had the other agent exerted an equal effort, and (ii) the other agent is assigned the residual output. The allocation under an RB rule is budget-balanced and independent of agents' costs of effort. We extend the RB rule to the case of n agents and consider its relationship to the common (but inefficient) rule of proportional allocation.

.

Keywords: Sharing rules, Reciprocity, Proportionality

JEL Classification Numbers: D24, C72, K13

< Preliminary >

1 Introduction

Suppose two agents worked on a project that turned out successful. Each agent independently performed a similar task, but one of them, call her the diligent agent, took more effort and her performance was consequently twice as good. How should the agents share the project's product? Or suppose that two agents engaged in a project that caused harm. Each agent conducted a similar activity, but one of them, call him the slack agent, took less care and therefore his activity was twice as risky. How should the agents share liability for the harm? The allocation of joint outputs among multiple agents, as in these examples, poses a twofold challenge. On the one hand, the allocation should possess basic fairness properties that would reflect agents' relative effort levels. On the other hand, the allocation should also provide agents with efficient incentives to exert efforts. In short, can allocation of outputs of joint efforts can be both fair and efficient?

A natural allocation, which is supposedly fair and efficient, is to divide the output (product or harm) in proportion to agents' efforts. Thus, in the examples above, the diligent agent's share in the product should be twice that of the slack agent, and the slack agent's liability share should be twice that of the diligent agent. But as Sen (1960) has shown in an influential paper, proportional allocation fails to provide efficient incentives. When returns to scale are non-constant, agents obtain (or are liable for) either more or less than their individual contribution to the joint output. Proportional allocation thus results in inefficient supply of efforts.

The failure of proportionality to induce optimal efforts might suggest that, in designing an optimal allocation, one must trade off fairness against efficiency. For example, Congress and the Supreme Court, guided by concern for fairness, have each resorted to proportionality as a principle of allocation.[1] In the private sector as well, fairness considerations seem to underlie the common practice of tying partners' shares in profits to their relative efforts (e.g., number of hours at work). This paper shows, however, that replacing *proportionality* with *reciprocity* as the standard for sharing yields both a fair and efficient outcome. In particular, we introduce an optimal allocation rule that is responsive to agents' relative efforts ex post and provides optimal effort incentives ex ante. Under this sharing rule, the slack agent obtains an equal share (one half) of the *hypothetical product* that would have transpired had the diligent agent been equally slack; the diligent agent obtains the residual of the *actual* product. The allocation of liability is similar: The diligent agent is liable for one half of the *hypothetical harm* that would have resulted had the slack agent been equally diligent; the slack agent is liable for the residual of the *actual* harm.

To study optimal allocation schemes, we consider two identical agents who engage in a joint project. Each agent must choose a level of effort with which to perform an activity, without observing the other agent's choice (this assumption is not crucial for our result). The higher the effort an agent takes, the more likely that its activity is

[1] Congress has advocated that liability of joint polluters be divided with particular reference to the relative "degree of care exercised by the parties" (American Jurisprudence §1355). The Supreme Court in *United States* v. *Reliable Transfer* established that, to achieve "just and equitable allocation," liability in admiralty law should be proportional to each injurer's carelessness level. Even the *Convention on International Liability for Damage Caused by Space Objects* provides that the liability of lunching states "shell be apportioned . . . in accordance with the extent to which they were at fault."

successful or safe. After the agents made their effort decisions, the output of the project materializes, depending on the agents' efforts. The allocation problem is to divide the output according to the agents' levels of effort so that agents fairly share the entire output (allocation is thus budget-balanced) and have optimal ex-ante incentives to exert effort. Because the problem of allocating harm is the mirror-image of that of allocating product, we concentrate throughout the paper mainly on the latter problem.

We begin by considering the different incentives transmitted by proportional allocation, whereby agents share the joint product in proportion to their efforts. Each agent has a direct incentive to exert effort, because he thereby obtains a (proportional) fraction of its contribution to the product. But proportional allocation also provides each agent a strategic incentive to exert effort, because exerting more effort increases an agent's relative share in the product at the expense of the other agent. The overall effect of proportional allocation on agents' incentives depends, as others have shown, on the properties of the production function (see, for example, Israelsen, 1980). If the production function has decreasing returns to scale, agents exert too much effort; if it has increasing returns to scale, agents exert too little effort.

Against this background, we present a simple allocation rule that induces agents to choose optimal effort levels irrespective of the properties of their production function. Under this rule, if agents exert equal efforts, each is assigned half of the product, as under proportional allocation. If agents' choices of effort differ, one agent is assigned one half of the output that would have obtained had the other agent exerted an equal effort. The other agent is assigned the residual of the actual output. Because one agent's share in the output is derived from hypothetical reciprocity, we call this allocation rule "Reciprocity Based" (RB).

The RB rule provides optimal incentives by aligning each agent's private benefit of effort (increase in product share, or decrease in liability share) with the corresponding social benefit (increase in product or decrease in harm). To see this, consider the RB rule whereby the slack agent's share in the product is based on hypothetical reciprocity. As we show in the paper, under this rule the slack agent's benefit of effort changes by the same sign as the change in the hypothetical social welfare given that both agents are equally slack. This implies that the slack agent has no incentive to take less effort than the socially-optimal one. The diligent agent's payoff is equal to social welfare less the slack agent's (fixed) payoff and therefore this agent's private benefit of effort is equal to the corresponding social benefit. The diligent agent consequently has no incentive to take more effort than the socially-optimal one. Thus, no agent can profitably deviate from the socially-optimal effort level by exerting either more or less effort. (The same reasoning applies to the RB rule whereby the diligent agent's liability is based on hypothetical reciprocity.)

To illustrate the force of reciprocity as a sharing principle, suppose that two agents engage in a joint project. Let e_i denote agent i's effort ($i = 1, 2$) and suppose that the project's product is $Q(e_1, e_2) = q(e_1) + q(e_2)$. Suppose further that, given the agents' costs of effort, there exists a unique pair of effort levels that maximizes social welfare. Now, under the RB rule, agent i obtains $q(e_i)$. In particular, if agent i's payoff is based on hypothetical reciprocity, his product share is $\frac{1}{2} \times (q(e_i) + q(e_i)) = q(e_i)$; agent j, who is entitled to the residual product, thus obtains $q(e_j)$. The same result holds if agent j's payoff is based on hypothetical reciprocity, and agent i is entitled to the residual product. Each agent thus obtains his exact contribution to the joint product.[2] Moreover, the allocation is *doubly* reciprocal in that each agent's payoff is equal to his payoff had the other agent exerted as much effort.

The allocative power of the RB rule is stronger yet when agents' efforts complement or substitute each other.[3] Each agent's contribution to the joint product in this case depends on the other agent's effort level. Because there is no clear way of disentangling agents' respective contributions to the product, proportional allocation is particularly appealing. However, here as well proportionality fails to provide efficient incentives. Reciprocity-based allocation, which prescribes an intuitive method for separating out agents' contributions, also provides optimal effort incentives.

Because the RB rule only requires information on the output function and the agents' efforts, it can be applied even in the absence of information on agents' costs of effort and therefore on the socially-optimal outcome (unlike the often-analyzed sharing rule that involves a mixture of equal and proportional shares; see, for example, Cornes and Hartley, 2004). The RB rule is also budget-balanced. The output share of the agent whose payoff is based on hypothetical reciprocity is less than one. Because the other agent is assigned the residual output, the entire output is allocated

The extension of the RB rule to multiple agents follows the contours of the two-agent case. In particular, an optimal allocation of product assigns each agent, other than the most diligent, $1/n$ of the hypothetical product that would have resulted had all the more diligent agents exerted as much effort; the residual product is shared equally among the most diligent agents (in case of a harm, the residual agents are the most slack). Each agent's payoff is thus reciprocal with respect to all the more diligent agents; in this sense it is "weakly reciprocal." Weak reciprocity is an upshot of the balanced-budget requirement when agents efforts have positive cross effects (i.e., if efforts complement each other). In the absence of complementarities in production, however, each agent other than the most diligent should obtain $1/n$ of the hypothetical product that would have resulted had all other agents exerted as much effort (the residual of the actual product is shared equally among the most diligent). This allocation thus satisfies "strong" reciprocity. If agents' efforts have no cross effect (as in the example above), *all* agents obtain their hypothetically reciprocal shares.

Finally, we show that the RB rule satisfies four intuitive fairness properties. First, an increase in an agent's effort increases its absolute share in profit and decreases its absolute share in liability (Monotonicity). Second, a more diligent agent obtains a higher profit share, or bears a lower liability share, than a less diligent agent (Equity). Third, equal agents are treated alike in that any two or more agents who exerted the same effort obtain the same profit or bear the same liability (Equality). Forth, any subset of agents (if $n > 2$) divide their aggregate output shares in accordance with the hypothetical reciprocity principle (Consistency).

This paper is part of a vast literature on optimal effort-based allocation rules. Our intention here is merely to provide a sample of this literature; for a comprehensive survey

[2]Under proportional allocation, each agent's (absolute) share in the product is $e_i/(e_1+e_2) \times Q(e_1, e_2)$, which is generally different from $q(e_i)$.

[3]Consider, for example, $Q(e_1, e_2) = k(e_1 + e_2)^{\frac{1}{2}}$, for $k > 0$.

see, for example, Bonin et al., 1993. To the best of our knowledge, however, no paper has examined the fairness and efficiency properties of reciprocity-based allocation rules.

One of the seminal papers on sharing rules is Sen (1966), who examined their incentive effects. Sen shows that proportional sharing creates excessive effort incentives when the production function exhibits decreasing returns to scale. To restore efficiency, Sen proposed a sharing rule under which a fraction of the output is distributed proportionally whereas the complementary fraction is shared equally. Sen's solution, unlike the one suggested here, requires information on the socially-optimal effort levels and is designed for production functions with decreasing returns to scale. Israelsen (1980) compares work incentives of communes (equal allocation), collectives (proportional allocation), and capitalistic firms (price allocation) and shows, similar to Sen, that the incentives to work in a collective are greater than in a firm or a commune under decreasing returns to scale. In the context of harm sharing, Kornhauser and Revesz (1989, 1995) show that proportional allocation of (strict) liability creates incentives to over-produce (i.e., exert too little care) if the harm function is convex, and therefore returns to scale are increasing. More recently, Friedman (2004) devised a sophisticated allocation technique using fixed-path methods, which satisfies a set of fairness axioms (for an economic interpretation see Leroux, 2007). Finally, Bachrach et al. (2012) provide a general characterization of the properties of (effort-based) distribution mechanisms that achieve a threshold level of efficiency.

The paper unfolds as follows. Section 2 presents the setup of the model and shows that proportional allocation distorts incentives if the production function displays non-constant returns to scale. Section 3 introduces the RB rule and shows that it implements the unique efficient Nash equilibrium, irrespective of the agents' production function. This Section then compares the allocation under the RB rule to proportional allocation and considers its fairness properties. Section 4 extends the RB rule to the case of multiple agents ($n > 2$). Section 5 concludes.

2 Model

2.1 Setup

Consider two risk-neutral agents, 1 and 2, who engage in a joint project. The project consists of two identical tasks, each performed by a different agent. The quality of each agent's performance on its task depends on the agent's effort; the more effort the agent exerts, the higher is the quality of its performance. We represent agent i's effort by $e_i \in [0, 1]$, where $i = 1, 2$. A convenient interpretation of e is that it represents the *probability* that the agent successfully completes its task. Under this interpretation, the ratio between the agents' efforts reflects the ratio between the probabilities that each successfully completes its task.

Each agent's cost of effort is $C(e)$, whish is strictly increasing and convex. We suppose that there is no cost to being idle and that the derivative of the cost with respect to effort tends to zero as e approaches 0 and to infinity as e approaches 1 ($C(0) = 0$, $C'_{e \to 0+}(e) = 0$

and $C'_{e\to 1-}(e) = \infty$). The function $C(e) = -k\,(e + \ln(1-e))$, which satisfies these conditions, will serve us in particular examples.[4]

Let $Q(e_1, e_2) : [0,1]^2 \to [0,1]$ be the probability that the project is successful as a function of the agents' efforts. We normalize the magnitude of the product from a successful project to 1 so Q also represents expected product. We assume that the project fails with certainty if both agents exert no effort and succeeds with certainty if both agents exert maximal effort ($Q(0,0) = 0$ and $Q(1,1) = 1$). We further assume that the probability of the project's success increases at a non-increasing rate with each agent's effort ($\frac{\partial Q}{\partial e_i} > 0$ and $\frac{\partial^2 Q}{(\partial e_i)^2} \leq 0$).

We will consider in particular three special cases. If effort are *complements*, the project succeeds iff both agents successfully complete their tasks (here we interpret effort as a probability of completing one's task). The probability of the project's success is therefore $e_1 e_2$. If efforts are *substitutes*, the project succeeds if either (or both) agents successfully completes its task. The probability of the project's success is accordingly $1-(1-e_1)(1-e_2)$. Finally, if efforts are *independent*, the probability of the project's success depends separately on each agent's effort and is given by $\frac{1}{2}e_1^\alpha + \frac{1}{2}e_2^\alpha$, where $\alpha \in (0,1]$. These cases represent, respectively, production functions with increasing, decreasing, and constant differences.[5]

Social welfare, which we assume is strictly quasi-concave, is equal to the expected product less the agents' costs of effort

$$Q(e_1, e_2) - C(e_1) - C(e_2). \tag{1}$$

To maximize social welfare, each agent's marginal cost of effort should be equal to the corresponding social marginal benefit (resulting from the higher expected product). Thus, at the socially-optimal outcome

$$C'(e_i) = \frac{\partial Q(e_1, e_2)}{\partial e_i} \text{ for } i = 1, 2. \tag{2}$$

We assume that there exists a unique (symmetric) efficient effort level and denote it by e^*.

2.2 Proportional Allocation

Under proportional allocation (PA), agent i's share in the product is $\frac{e_i}{E}$, where $E = e_1 + e_2$. Each agent's share thus reflects its *proportional* effort. Agent i's maximization problem under PA is therefore

$$\max_{e_i} \frac{e_i}{E} Q(e_1, e_2) - C(e_i). \tag{3}$$

[4] Note that $C'(p) = \frac{kp}{1-p} > 0$ and $C''(p) = \frac{k}{(1-p)^2} > 0$.

[5] A function $Q(e_1, e_2)$ has increasing (decreasing) differences if for every $\underline{e}_1, \overline{e}_1$ and $\underline{e}_2, \overline{e}_2$ such that $\overline{e}_2 > \underline{e}_2$ and $\overline{e}_1 > \underline{e}_1$ we have $Q(\overline{e}_1, \overline{e}_2) - Q(\underline{e}_1, \overline{e}_2) > (<) Q(\overline{e}_1, \underline{e}_2) - Q(\underline{e}_1, \underline{e}_2)$. This means that the incremental increase in the product from a higher effort level of one agent increases (decreases) if the other agent exerts more effort.

Differentiating with respect to e_i, equating to zero and rearranging yields

$$C'(e_i) = \frac{e_i}{E}\frac{\partial Q}{\partial e_i} + \frac{e_j}{E}\frac{Q}{E}, \tag{4}$$

where the arguments of Q are suppressed. The left-hand side is the agent's marginal cost of effort. The right-hand side is the corresponding marginal benefit. Each agent's benefit of effort consists of direct and strategic benefits (first and second term on the RHS of (4), respectively). The *direct* benefit stems from the increase in the expected product times the agent's share in the product. Because each agent obtains a (proportional) fraction of the total product, exerting more effort confers a positive externality on the other agent (note that the *social* benefit of effort is $\partial Q/\partial e_i$). The *strategic* benefit stems from the redistribution of product shares that results when one agent increases its effort. As one agent takes more effort, it increases relative product share at the expense of the other agent. Exerting more effort thus imposes a negative externality on the other agent.

The next Proposition compares the agents' private choices of effort and the socially-optimal ones.

Proposition 1 *Suppose agents share the product proportionally. Then:*

(a) There exists a unique symmetric equilibrium of effort levels.

(b) Agents exert too little (much) effort if, at the socially-optimal outcome, Q exhibits increasing (decreasing) returns to scale.

Proof. See the Appendix.

The intuition for the inefficiency of PA follows from the fact that each agent's marginal benefit of effort (RHS of (4)) is a convex combination of the marginal product of effort ($\partial Q/\partial e_i$) and the average product of effort (Q/E) (see Israelsen, 1980). Whether agents exert too little or too much effort thus depends on the relation between the marginal and the average product of effort. If the marginal product is equal to the average product (constant returns to scale), each agent's strategic benefit of effort is equal (in equilibrium) to one half of the marginal product of effort. As a result, agents' incentives are aligned with social welfare. The case of independent efforts with $\alpha = 1$ represent constant returns to scale ($Q = (e_1 + e_2)/2$). If the marginal product is greater or smaller than the average product (increasing or decreasing returns to scale), agents exert too little or too much effort, respectively. The case of independent efforts with $\alpha < 1$ represent decreasing returns to scale ($Q = (e_1^\alpha + e_1^\alpha)/2$).[6]

Figure 1 shows each agent's socially-optimal effort level (bold curves) as well as privately-optimal effort level for complementary and substitute efforts under PA (dashed curves), for each choice of effort of the other agent.[7] It also presents iso-welfare curves for

[6] This follows because $\partial Q/\partial e_i = \frac{1}{2}\alpha\frac{e_1^\alpha}{e_1} < \frac{1}{2}\frac{e_1^\alpha}{e_1} < \frac{1}{2}\frac{e_1^\alpha + e_2^\alpha}{e_2+e_2} = Q/E$ for $\alpha < 1$.

[7] Note that agent 1's *socially-optimal* reaction curve passes through the uppermost (lowermost) points of the iso-welfare curves for $e_1 > (<)\ e_2$ and that agent 2's *socially-optimal* reaction curve passes through the leftmost (rightmost) points of the iso-welfare curves for $e_2 > (<)\ e_1$. At these points, the marginal cost of effort is equal to the corresponding marginal social benefit (from the higher expected product).

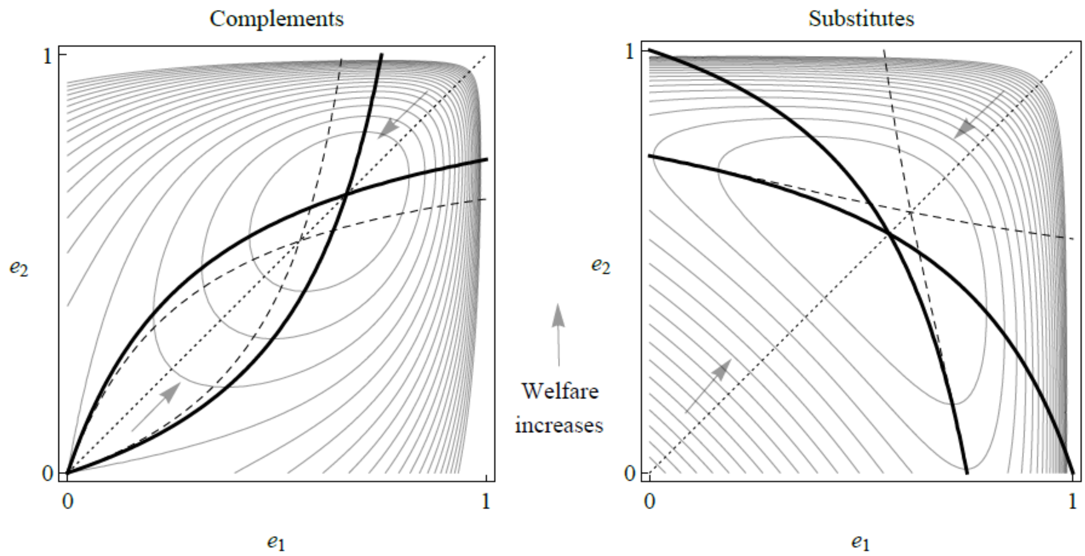

Figure 1: Social and Private Reaction Curves (under proportional allocation)

different choices of effort of each agent. The *socially-optimal* choice of efforts is located at the intersection of the socially-optimal best-response curves, whereas the agents' *equilibrium* choice of efforts is attained at the intersection of the agents' privately-optimal best-response curves. This figure thus shows the divergence between the social and private incentives to exert effort under PA for a production function with positive or negative *cross effects* (increasing or decreasing differences). In particular, if efforts are complements, returns to scale are increasing and agents exert insufficient efforts relative to the socially-optimal efforts. If efforts are substitutes, returns to scale are decreasing and agents exert excessive efforts.[8]

3 Reciprocity-Based Allocation

This section introduces a balanced-budget sharing rule that implements the socially-optimal effort levels irrespective of the properties of the production function. Call "diligent" the agent who exerted more effort and "slack" the agent who exerted less effort. If both agents exerted the same effort, suppose both are diligent. We will use the term "product" to refer to the *expected* product, "product share" to refer to each agent's expected *absolute* product share, and "relative product share" to refer to each agent's *percentage* share in the product.

Under the Reciprocity-based (RB) rule: (i) the slack obtains one half of the hypothetical product that would have resulted had the diligent been equally slack; and (2) the diligent obtains the residual of the product. We state the RB rule formally in the following definition.

[8]Note that in the complements case, $\partial Q/\partial e_i = e_j > \frac{e_1 e_2}{e_1+e_2} = Q/E$ (increasing returns to scale), and in the substitutes case $\partial Q/\partial e_i = 1 - e_j < 1 - \frac{e_1 e_2}{e_1+e_2} = Q/E$ (decreasing returns to scale).

Definition 1 (Reciprocity-Based Sharing Rule) *Under the RB rule, agent i's relative product share as a function of its and agent j's effort is*

$$\begin{cases} \frac{1}{2}Q(e_i,e_i)/Q(e_1,e_2) & \text{if } e_i < e_j \\ 1-\frac{1}{2}Q(e_j,e_j)/Q(e_1,e_2) & \text{if } e_i \geq e_j \end{cases}.$$

To see that the RB rule is balanced-budget, let e_d and e_s denote the diligent's and the slack's effort levels, where $e_d \geq e_s$ (in case of an equality, suppose both agents are diligent). Call $Q(e_s,e_s)$ the "reciprocal product"; that is, the product that would have resulted had both agents been equally slack. Note that, because the product increases with each agent's effort, the reciprocal product is lower than the product ($Q(e_s,e_s) < Q(e_1,e_2)$). Now, the slack's relative product share is equal to one half of the reciprocal product divided by the product and therefore is strictly less than 1. Because the diligent obtains the complementary share, the RB rule balances the budget.

Before considering the equilibrium under the RB rule, it would be helpful to formally specify each player's expected payoff under the RB rule. The slack's payoff is one half of the reciprocal product less its own cost of effort: $\frac{1}{2}Q(e_s,e_s) - C(e_s)$. The diligent's payoff is equal to the product less its own cost of effort and the slack's payoff: $Q(e_1,e_2) - \frac{1}{2}Q(e_s,e_s) - C(e_d)$. The next proposition now turns to the efficiency of the RB rule.

Proposition 2 *The unique Nash equilibrium under the RB rule is for each agent to exert the socially-optimal effort.*

Proof. The proof proceeds in three steps. We first show that, given that both agents choose the socially-optimal effort level, e^*, no agent can profitably deviate to taking either more or less effort. We then show that there is no equilibrium in which both agents choose the same suboptimal effort level. Finally, we show that there is no equilibrium in which the agents choose different effort levels.

Suppose that both agents choose the socially-optimal effort level, e^*. Let $W(e_1,e_2) = Q(e_1,e_2) - C(e_1) - C(e_2)$ denote social welfare as a function of the agents' efforts. Observe that the slack's payoff is equal to one half of social welfare along the 45-degree line: $\frac{1}{2}Q(e_s,e_s) - C(e_s) = \frac{1}{2} \times W(e_s,e_s)$. In contemplating a downward deviation, an agent must therefore consider the change in social welfare along the 45-degree line. Because the gradient $(\frac{\partial W}{\partial e_1}, \frac{\partial W}{\partial e_2})$ at any point (e,e), where $e < e^*$, is positive (from the FOC in (2); observe the iso-welfare curves in Figure 1), social welfare increases up along the 45-degree line. A deviation to a *lower* effort level, which reduces the hypothetical social welfare given that both agents exert equal effort, is therefore unprofitable.

Next, observe that the diligent's payoff, $Q(e_1,e_2) - \frac{1}{2}Q(e_s,e_s) - C(p_d)$, is equal to social welfare less the slack's fixed payoff: $W(e_1,e_2) - [\frac{1}{2}Q(e_s,e_s) + C(p_s)]$. In contemplating an upward deviation, an agent must therefore consider the change in social welfare induced by exerting more effort, given that the other agent chooses the socially-optimal effort. That is, the agent must consider the change in social welfare along a vertical or horizontal shift (depending on whether agent 1 or 2 deviates) from the socially-optimal effort levels. But from the FOC in (2), the socially-optimal best-response to the socially-optimal effort level is to respond in kind (by choosing e^*). Thus a deviation to a *higher* effort level is unprofitable.

Consider next a putative *symmetric* equilibrium where both agents choose the same suboptimal effort $e \neq e^*$. We proceed by considering two cases. In the first case, both agents choose an inefficiently *high* effort level $\overline{e} > e^*$. In the second case, both agents choose inefficiently *low* effort level $\underline{e} < e^*$. Suppose first that both agents choose $\overline{e} > e^*$. Each agent's equilibrium payoff is therefore $\frac{1}{2}W(\overline{e})$. But any agent can profitably deviate down to e^*, the socially-optimal level (thereby becoming slack). To see this, note that the sign of the change in the slack's payoff from taking less effort is opposite to the sign of the gradient $(\frac{\partial W}{\partial e_1}, \frac{\partial W}{\partial e_2})$ along the 45-degree line. But the sign of the gradient at any point (e, e), where $e > e^*$, is negative (from the FOC in (2)). If follows that $\frac{1}{2}W(e^*) > \frac{1}{2}W(\overline{e})$, and therefore a deviation to taking the socially-optimal effort e^* is profitable.

Suppose next that both agents choose an effort level $\underline{e} < e^*$. Then any agent can profitably deviate to taking more effort (thereby becoming the diligent). To see this, consider (w.l.o.g.) a deviation of agent 1 to taking an effort level $\widehat{e}$, where $\widehat{e} > \underline{e}$, such that social welfare given $(\widehat{e}, \underline{e})$ is identical to social welfare given $(\underline{e}, \underline{e})$; that is $W(\widehat{e}, \underline{e}) = W(\underline{e}, \underline{e})$. Such $\widehat{e}$ exists because the direction of the gradient $(\frac{\partial w}{\partial e_1}, \frac{\partial w}{\partial e_2})$ at a point along the 45-degree line is perpendicular to the tangent line through such a point. This in turn implies that the iso-welfare curve through a point $(\underline{e}, \underline{e})$, where $\underline{e} < e^*$, slopes downward with a slope of $\frac{3}{4}\pi$. By strict quasi-concavity of W, social welfare of any convex combination of $(\widehat{e}, \underline{e})$ and $(\underline{e}, \underline{e})$ is greater than social welfare at $(\underline{e}, \underline{e})$: $W[\lambda(\widehat{e}, \underline{e}) + (1 - \lambda)(\underline{e}, \underline{e})] > W(\underline{e}, \underline{e})$, where $\lambda \in (0, 1)$. This implies that the privately- (as well as socially-) optimal choice of e_1, given $e_2 < e^*$, is strictly greater than e_2, and therefore a deviation to an effort level that maximizes social welfare is profitable.

Finally, consider a putative *asymmetric* equilibrium, where one agent chooses a higher effort level than the other. Then the slack can profitably deviate to taking more effort if its initial effort level is lower than e^* (because social welfare increases up along the diagonal for effort levels lower than e^*); and to taking less effort than if its initial effort level is greater than e^*(because social welfare decrease up along the diagonal for effort levels higher than e^*). ■

The RB rule implements the socially-optimal effort levels by aligning each agent's payoff with social welfare. In particular, the change in the diligent's payoff from exerting more effort is *equal* to the corresponding change in social welfare. This is because the diligent's payoff is equal to social welfare less the slack's fixed payoff (expected product share plus cost of effort). The diligent's payoff, consequently, is independent of the slack's choice of effort. The diligent thus fully internalizes the social benefit (or cost) of taking more effort.

The slack's payoff from exerting more effort changes *by the same sign* as the change in social welfare along the 45-degree line. This is because the slack's payoff equals one half of social welfare given that both agents exert equal effort. The slack thus has incentives to increase effort up to the socially-optimal effort level (given that the diligent exerts more effort), but never more than the socially-optimal effort.

Moreover, in the absence of cross effects between the agents' efforts (Q has constant differences), the change in the slack's payoff following an increase in its effort is *equal* to the corresponding change in social welfare; it thus does not affect the other (diligent) agent's payoff. To see this, suppose (w.l.o.g.) that agent 1 is slack. Then a marginal increase in agent 1's effort changes social welfare by $\frac{\partial Q(e_1,e_2)}{\partial e_1} - C'(e_1)$. The corresponding change in agent 1's payoff is $\frac{1}{2}\frac{dQ(e_1,e_1)}{de_1} - C'(e_1)$. But the first term is equal to the partial derivative of Q with respect to e_1 given that $e_1 = e_2$ $\left(\left.\frac{\partial Q(e_1,e_2)}{\partial e_1}\right|_{e_2=e_1}\right)$. If Q has no cross effects, this partial derivative is independent of agent 2's effort $\left(\left.\frac{\partial Q(e_1,e_2)}{\partial e_1}\right|_{e_2=e_1} = \left.\frac{\partial Q(e_1,e_2)}{\partial e_1}\right|_{e_2\neq e_1}\right)$. Thus, in the absence of cross effects, the change in the slack's payoff from taking more effort is equal to the corresponding change in social welfare.

In the presence of cross effects between the agents' efforts, an increase in the slack's effort affects the diligent's payoff. A higher effort of the slack changes that agent's payoff by either more or less than the corresponding change in social welfare. In particular, if Q has decreasing differences (negative cross effects), an increase in the slack's effort imposes a negative externality on the diligent (whose payoff decreases); if Q has increasing differences (negative cross effects), an increase in the slack effort confers a positive externality on the diligent (whose payoff increases). Irrespective of the properties of Q, however, the marginal change in the slack's payoff from a higher effort converges to the corresponding change in social welfare as the slack's effort approaches the socially-optimal one. This means that, given that one agent exerts the socially-optimal effort, the other agent's marginal deviation down from the socially-optimal effort reduces the deviating agent's payoff by exactly the corresponding reduction in social welfare.

Figure 2 presents the agents' (private) best-response curves under the RB rule. This figure shows that if the other agent chooses an effort level higher than the socially-optimal level, each agent's best response is to choose the socially-optimal effort level (e^*). If the other agent chooses an effort level lower than the socially-optimal level, each agent's best response is to choose the socially-optimal best response, which is *greater* than the other agent's choice of effort. More specifically, the socially (as well as privately) optimal choice of effort, given that one agent's effort is lower than the socially-optimal effort, is greater than e^* if Q has increasing differences and lower than e^* if Q has decreasing differences. If Q has constant differences, each agent's best response to *any* choice of effort of the other agent is to apply the socially-optimal effort level, e^*.

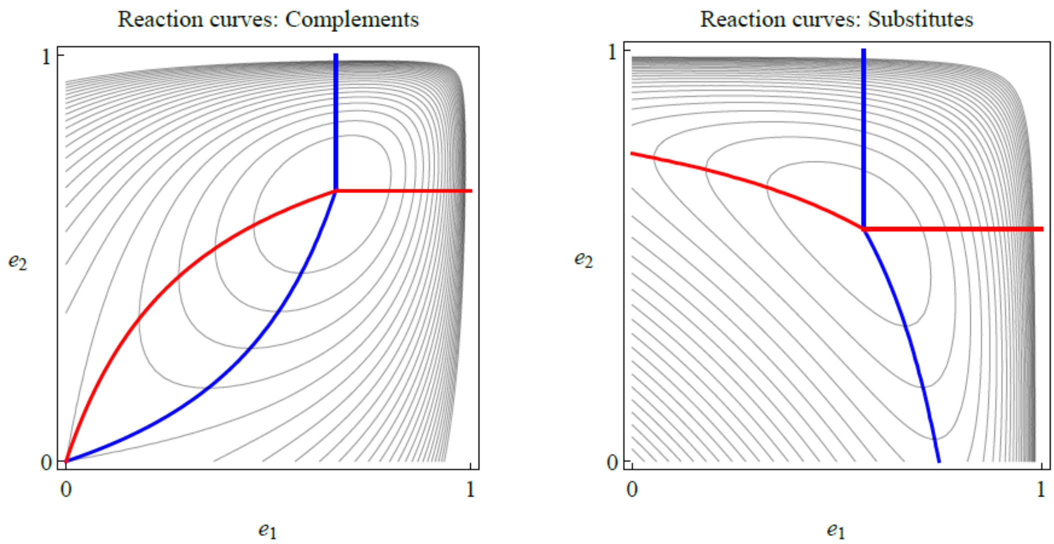

Figure 2: Reaction Curves under the Reciprocity-Based Rule

We now turn to comparing the RB rule and a proportional rule. Recall that under PA an agent who marginally deviates from the socially-optimal outcome obtains more (less) than its corresponding contribution to the product if the production function has decreasing (increasing) returns to scale. Under the RB rule, in contrast, the increase in such an agent's payoff is exactly equal to its contribution to the product. This means that in the vicinity of the socially-optimal outcome, the RB rule grants the marginally more diligent agent more (less) than its proportional share if the production function has decreasing (increasing) returns to scale. The next proposition more generally considers the relationship between proportional and reciprocal allocations.

Proposition 3 *The diligent's relative product share under a reciprocal allocation is lower than its proportional share if Q has decreasing returns to scale and non-decreasing differences and is greater than its proportional share if Q has increasing returns to scale and non-increasing differences.*

Proof. See the Appendix.

The allocation under the RB rule depends on both scale effects and cross effects. In the absence of cross effects, an increase in one agent's effort increase the product by the *average* increase in the product when both agents increase their efforts from e_s to e_d ($\overline{Q}(e_i, e_j) := \frac{1}{2}[Q(e_d, e_d) - Q(e_s, e_s)]$). Under the RB rule, therefore, the increase in the diligent's payoff is equal to the the increase in the product (and is therefore equal to $\overline{Q}(e_i, e_j)$). Under PA, in contrast, a diligent agent obtains more than its contribution to the increase in the product if returns to scale are decreasing and less than its contribution to the product if returns to scale are increasing (hence efforts are excessive in the first case and insufficient in the second). A reciprocal allocation thus grants the diligent agent less than its proportional share if returns to scale are decreasing and more than its proportional share if returns to scale are increasing (vice versa for the slack).

The presence of cross effects determines whether an increase in one agent's effort increase the *product* by more or less than $\overline{Q}(e_i, e_j)$. In particular, in the presence of positive (negative) cross effects, an increase in one agent's effort increases the product by less (more) than $\overline{Q}(e_i, e_j)$. It follows that if Q has decreasing returns to scale and positive cross effect (non-decreasing differences), the increase in the product that results from an increase in one agent's effort is less than $\overline{Q}(e_i, e_j)$. As a result, a reciprocal allocation is tilted in favor of the slack (because the diligent obtains strictly less than $\overline{Q}(e_i, e_j)$ under reciprocal allocation but more than than $\overline{Q}(e_i, e_j)$ under PA). By a similar reasoning, if Q has increasing returns to scale and negative cross effects (non-increasing differences), a reciprocal allocation is tilted in favor of the diligent.[9]

Interestingly, the RB rule implements proportional allocation if Q has constant returns to scale and no cross effects. In this case, given the diligent's effort is β times the slack's effort ($\beta > 1$), the product increases $\frac{1+\beta}{2}$ times. The slack's proportional share is thus equal to his reciprocal share (because $\frac{1}{1+\beta} = \frac{1}{2}Q(e_s, e_s)/\left(\frac{1+\beta}{2}Q(e_s, e_s)\right)$); hence the diligent's proportional share is equal to his reciprocal share as well. Proportional

[9]To see the interaction between scale effects and cross effects, consider the following Cobb-Douglas production function: $Q(e_1, e_2) = (e_1 e_1)^{\frac{1}{2}}$. It is easy to verify that this production function has constant returns to scale as well as positive cross effects.

allocation is thus a special case of reciprocal allocation in the case in which the production function has constant returns to scale and no cross effects.[10]

We conclude this section by considering four fairness properties of the RB rule. First, the RB rule is *egalitarian*, because agents share the product equally if they choose the same effort level. Second, the RB rule is *monotonic*, because each agent's product share increases with its effort. The diligent's product share increases with its effort because this agent obtains the residual of the product less the slack's fixed share. A higher effort, which increases the product, thus results in a higher absolute as well as relative product share. The slack's product share increases with its effort because the reciprocal product (given that both agents are equally slack) increases as both agents exert more efforts. Moreover, if Q has non-increasing differences (i.e., negative or no cross effect), the slack's *relative* product share increases with its effort as well. This follows because in this case an increase in the slack's effort decreases the diligent's payoff, as we showed above. Because the increase in the slack's product share is greater than the corresponding increase in the product, the slack's relative product share must increase. Third, the RB rule is *equitable*, because the diligent's product share is greater than the slack's share. To see this, observe that the product is greater than the hypothetical reciprocal product ($Q(e_1, e_2) \geq Q(e_s, e_s)$), which in turn implies that the diligent's share is greater than the slack's share ($Q(e_1, e_2) - \frac{1}{2}Q(p_s, p_s) \geq \frac{1}{2}Q(p_s, p_s)$, with strict inequality for $e_d > e_s$).

Finally, we consider an additional fairness property, which we call *reciprocity.* Reciprocity requires that each agent's payoff be greater or equal to its payoff had the other agent exerted equal effort. We now show that if the production function has no cross effects or negative cross effects (that is, if the production function has non-increasing differences) the RB rule is *fully reciprocal* in that each agent's payoff is equal to or is greater than its reciprocal payoff. Moreover, in the absence of cross effects (Q has constant differences), the RB rule is *doubly reciprocal* in that each agent's payoff is *equal* to its payoff had the other agent exerted equal effort.

To see this, note that if Q has non-increasing differences, then $Q(e_d, e_d) - Q(e_s, e_d) \leq Q(e_d, e_s) - Q(e_s, e_s)$. This implies that, given that one agent is more diligent than the other, the sum of the agents' reciprocal product shares is lower than the product: $\frac{Q(e_d,e_d)+Q(e_s,e_s)}{2} \leq Q(e_s, e_d)$ (because $Q(e_d, e_s) = Q(e_s, e_d)$). But this means that the diligent's payoff under the RB rule is greater than its reciprocal payoff: $\frac{Q(e_d,e_d)}{2} - C(e_d) \leq Q(e_s, e_d) - \frac{Q(e_s,e_s)}{2} - C(e_d)$. Moreover, if Q has constant differences, then the last inequality becomes an equality and therefore both the diligent's payoff and the slack's payoff under the RB rule are *equal* to their reciprocal payoffs.

Finally, if Q has non-increasing differences, an RB rule under which the diligent's product share is reciprocal is balanced-budget as well. Under such rule, the diligent's payoff is equal to its reciprocal payoff and the slack's payoff is equal to or is greater than its reciprocal payoff. This rule, however, violates a desirable property of an allocation rule: a slack agent may obtain a positive share in the product even if it exerted no effort.

[10] If returns to scale are constant, a reciprocal allocation is tilted in favor of the slack (relative to PA) if there are positive cross effects and in favor of the diligent if there are negative cross effects.

4 N-Agents Reciprocity-Based Allocation

This section extends the RB rule to more than 2 agents; we call this rule Multi-Party Reciprocity-Based (MRB) rule. Consider a set $N = \{1, ..., n\}$ of agents, who must each choose a level of effort. Each agent has an identical cost of effort, which is increasing and convex and satisfies the Inada conditions set forth in the previous section. The probability of the project's success as a function of the agents' efforts is $Q(e_1, ..., e_n) : [0,1]^n \to [0,1]$. As in the previous section, we suppose that $\frac{\partial Q}{\partial e_i} > 0$ so that the expected product increases with each agent's effort. We further suppose that the social welfare function is strictly quasi-concave and that there exist a unique vector of symmetric effort levels, $(e^*, ..., e^*)$, that maximizes social welfare.

As in the case of $n = 2$, a proportional allocation distorts agents' incentives to take effort. In particular, agents exert too little effort if Q has increasing returns to scale, and too much effort if Q has decreasing returns to scale. The proof follows the same lines of the $n = 2$ case.

Under the MRB rule, each agent's product share, other than the most diligent agents, is equal to $1/n$ of the product that would have resulted had all the *more diligent* agents exerted as much (and therefore lower) effort. The most diligent agents share the residual product in equal shares. A formal definition follows.

Definition 2 (Multi-Party Reciprocity-Based Sharing Rule) *For each agent $i = 1, ..., n$, let $S_i = \{j \in N : e_j < e_i\}$ be the set of agents who exerted (strictly) less effort and $D_i = \{j \in N : e_j \geq e_i\}$ be the set of agents who exerted equal or more effort. For each agent $i = 1, ..., n$ and agent $j = 1, ..., n$, let $\widehat{e}_{ji} = e_i$ if $e_j \in S_i$ (i.e., if agent j exerted less effort than agent i) and $\widehat{e}_{ji} = e_i$ if $e_j \in D_i$ (i.e., if agent j exerted equal or greater effort than agent i). Then, letting Q denote the product given agents' efforts ($Q = Q(e_1, ..., e_n)$), agent i's relative product share under the MRB rule is*

$$\begin{cases} \frac{1}{n}Q(\widehat{e}_{1i}, ..., \widehat{e}_{ni})/Q & \text{if} \quad e_i < \max\{e_1, ..., e_n\} \\ \frac{1}{|D_i|}[1 - \sum\limits_{k \in S_i} \frac{1}{n}Q(\widehat{e}_{1k}, ..., \widehat{e}_{nk})/Q] & \text{if} \quad e_i = \max\{e_1, ..., e_n\} \end{cases}.$$

To see that the MRB rule is balanced-budget, consider the sum of the relative product shares of all agents other than the most diligent. Note that for any agent other than the most diligent (that is, $\{i \in N : e_i < \max\{e_1, ..., e_n\}\}$, we have $\widehat{e}_{ji} \leq e_j$ for $j < i$ and $\widehat{e}_{ji} = e_j$ for $j \geq i$, and therefore that $\frac{1}{n}Q(e_1, ..., e_n) > \frac{1}{n}Q(\widehat{e}_{1i}, ..., \widehat{e}_{ni})$. The equal share $(1/n)$ of each such agent in the hypothetical reciprocal product (RHS) is lower than its equal share in the product (LHS). Because there are at most $n-1$ sets of agents who exert equal efforts, the sum of relative product shares of agents who exerted less effort than the most diligent is less than 1: $\sum\limits_{k \in S_i} \frac{1}{n}Q(\widehat{e}_{1k}, ..., \widehat{e}_{nk})/Q \leq \frac{n-1}{n}Q(\widehat{e}_{1k}, ..., \widehat{e}_{nk})/Q < 1$, where i is the lowest-indexed agent who exerted the highest effort. Thus, the entire product is allocated among the n agents.

We now show that, as in the case of $n = 2$, the MRB rule induces socially-optimal efforts.

Proposition 4 *The unique strong Nash equilibrium under the MRB rule is for each agent to exert the socially-optimal effort.*

Proof. Suppose that each agent applies the socially-optimal effort, e^*. Then any unilateral deviation to a *lower* effort level yields the deviating agent $1/n$ of social welfare given that all other agents choose the same effort level. But social welfare obtains a unique maximum at $(e^*, ..., e^*)$, and therefore such a deviation is unprofitable. Next consider a unilateral deviation to a *higher* effort level. The (now) most diligent agent's payoff is equal to social welfare less all other agents' product shares plus costs of efforts. But the deviation decreases social welfare and therefore must decrease the deviating agent's payoff.

We now show that the equilibrium under the MRB rule is resilient to *multi-lateral* deviation (strong equilibrium). The solution concept of strong equilibrium requires that no coalition of agents could jointly deviate from their equilibrium strategies such that each member of the coalition is better off (unlike the coalition-proof refinement, the notion of strong equilibrium does not require that the joint deviation itself be self-enforcing). Consider then a coalition $R \subseteq N$ of r agents, where $r \in \{2, ...n\}$. Suppose that each agent i in R deviates to $e_i \neq e^*$ (note that more than one agent in R may deviate to the same effort level). Now, any agent in R who deviates to $e < e^*$ is worse off because any such agent's payoff is strictly lower than $1/n$ of social welfare given that all agents (both more diligent and more slack) exert as much effort. To see why, note that for any agent i, $\widehat{e}_{ji} \leq e_i$ for $j < i$ and $\widehat{e}_{ji} = e_i$ for $j \geq i$ and therefore $\frac{1}{n}Q(e_i, ..., e_i) > \frac{1}{n}Q(\widehat{e}_{1i}, ..., \widehat{e}_{ni})$. Each agent in R who deviated down thus obtains less than $\frac{1}{n}Q(e_i, ..., e_i) - C(e_i)$, which is equal to $\frac{1}{n}W(e_i, ..., e_i)$. But social welfare obtains a unique maximum at $(e^*, ..., e^*)$ and therefore such agent's payoff must be lower following the deviation. This means that any agent in R must deviate to taking more effort than the socially-optimal one. But social welfare decreases following the joint deviation and the payoff of any non-deviating agent's remains the same (because a slack agent's payoff is unaffected by another agent's deviation to a higher effort level if that agent initially took equally or more effort). It follows that the sum of payoffs of all agents in R must decrease and therefore that the joint deviation cannot make all members in R better off. We thus conclude that the equilibrium under the MRB rule is strong.

Finally, by arguments similar to the case of $n = 2$, there does not exist any other symmetric equilibrium or an asymmetric equilibrium. ■

As in the case of $n = 2$, the MRB is egalitarian, monotonic, and equitable. Egalitarianism follows because any two agents who exert equal efforts are treated alike. Monotonicity is satisfied because an agent's increase of effort increases its *absolute* product share. Moreover, if Q has non-increasing differences, an increase in an agent's effort also increases its *relative* product share (because, as in the case of $n = 2$, it imposes a negative externality on all the more diligent agents).

We now show the MRB is *equitable* in the sense that the product share of a more diligent agent is greater than that of a more slack agent. First note that any agent other than the most diligent obtains an equal share $(1/n)$ of the hypothetical product and that that hypothetical product monotonically increases with the agent's effort. Thus, an agent who exerts more effort (other than the most diligent) obtains a greater share in the product. Next, let $e_{\max} = \max\{e_1, ..., e_n\}$ denote the highest effort exerted by one or more agents. Suppose that $e_i < e_{\max}$ for some $i \in N$ so at least one agent exerted less effort than other agents. Let $D^1 = \{i \in N : e_i = e_{\max}\}$ be the set of agents who exerted the most effort and $D^2 = \{i \in N : e_i \geq e_k, k \in N/D^1\}$ the set of the second-most diligent agents. We now show that the product share of any agent in D^1 (most diligent) is lower than that of any agent in D^2 (second-most diligent). This implies that the product share of any agent who exerted the most effort is greater than the product share of any other agent.

To see this, denote the hypothetical product of any agent in D^2 (second-most slack) by $\widehat{Q}$; thus $\widehat{Q} := Q(\widehat{e}_{1k}, ..., \widehat{e}_{nk})$, where $k \in D^2$. Let $d = |D_1|$ denote the number of agents in D^1. Then, because $Q > \widehat{Q}$, it follows that $nQ - n\widehat{Q} + d\widehat{Q} > d\widehat{Q}$. Dividing through by dn gives $\frac{1}{d}\left(Q - \frac{\widehat{Q}(n-d)}{n}\right) > \frac{\widehat{Q}}{n}$. But the left-hand side is equal to or is greater than the product share of an agent in D^1 (most diligent), whereas the right-hand side is equal to the product share of an agent in D^2 (second-most-diligent). An agent who exerted the highest effort thus obtains a greater product share than an agent who exerted less effort.

We now consider reciprocity. The MRB rule satisfies *weak reciprocity* in that each agent other than the most diligent obtains an equal share of the hypothetical product that would have resulted had all the *more diligent* agents exerted equal effort. The most slack agent obtains its reciprocal payoff (i.e., its payoff had all other agents been equally slack). Moreover, if Q has non-increasing difference, the most diligent agent's payoff is higher than its hypothetical reciprocal payoff (its payoff had all other agents had been equally diligent). This follows because, for Q with non-increasing differences, $\frac{1}{n}\Sigma_{i=1}^{n} Q(e_i, ..., e_i) \leq Q(e_1, ..., e_n)$. This means that the sum of all agents' reciprocal product shares is lower than the product. Because each agent obtains no more than its reciprocal product share under the MRB rule, the most diligent agent must obtain more than its reciprocal product share.

We now show that if Q has non-increasing differences, then a *fully-reciprocal* MRB rule implements the efficient strong Nash equilibrium. Under a fully-reciprocal MRB rule every agent other than the most diligent obtains $1/n$ of the hypothetical product that would have resulted had all other agents exerted as much effort. The most diligent agents (or agent) obtains the residual product. Formally, under a *fully-reciprocal* MRB rule agent i's share in the product is

$$\begin{cases} \frac{1}{n}Q(e_i, ..., e_i)/Q & \text{if} \quad e_i < \max\{e_1, ..., e_n\} \\ \frac{1}{|D_i|}[1 - \sum\limits_{k \in S_i} \frac{1}{n}Q(e_i, ..., e_i)/Q] & \text{if} \quad e_i = \max\{e_1, ..., e_n\} \end{cases}.$$

A fully-reciprocal MRB rule is balanced-budget because the sum of all agents' reciprocal product shares is lower than the actual product. As a result, the entire product can be allocated by either (i) granting each agent, other than the most diligent, its reciprocal product share and granting the most diligent the residual product; or (ii) granting each agent, other than the most slack, its reciprocal product share and granting the most slack the residual product. Such a reciprocal rule thus satisfies *strong reciprocity* in that $n-1$ agents' payoffs are equal to their reciprocal payoff and one agent's payoff (either the most diligent or the most slack) obtains at least its reciprocal payoff. As in the case of $n = 2$, however, an MRB sharing rule under which the diligent's payoff is based on hypothetical reciprocity may grant a slack agent a positive product share even if it exerted no effort.

Before we conclude we consider the *consistency* of the MRB rule (see Aumann and Maschler, 1985). Specifically, consider any subset $R \subset N$ of agents. Let Q_R denote the aggregate product shares of agents in R. Consistency requires that the allocation of Q_R among all agents in R satisfies the MRB with respect to the original production function, $Q(e_1, ..., e_n)$. Consistency is satisfied because the product share of any agent in R, other than the most diligent, is based on hypothetical reciprocity, whereas the most diligent agents in R obtain the residual.

5 Conclusion

This paper proposed an effort-based mechanism for allocating joint outputs among identical agents. The challenge in designing an optimal allocation scheme is in satisfying basic fairness properties and providing optimal effort incentives. We showed that reciprocity is a powerful principle for achieving both goals. In the two-agent case, a reciprocity-based allocation treats one agent as if both agents exerted as much effort and assigns the residual output to the other agent. It thus provides an intuitive way of separating out each agent's contribution to the joint output. Equally important, a reciprocity-based allocation aligns each agent's payoff with its corresponding contribution to social welfare.

An interesting extension is the case of asymmetric agents, where agents have different costs of effort (or different contributions to the output). Unlike the case in which agents are identical, it is not clear how to implement a reciprocity-based allocation. We hope to take up this question in another paper.

Appendix

This Appendix restates and proves Propositions 1 and 3.

Proposition 1 *Suppose agents share the product proportionally. Then:*

(a) There exists a unique symmetric equilibrium of effort levels.

(b) Agents exert too little (much) effort if, at the socially-optimal outcome, Q exhibits increasing (decreasing) returns to scale.

Proof. (a) Suppose that each agent chooses an effort level $\widehat{e}$. Note that the LHS of (4) increases with $\widehat{e}$, because the agents' cost function is convex. Plugging $\widehat{e}$ for e_i in the the RHS of (4) yields $\frac{1}{2}\frac{\partial Q(e,\widehat{p})}{\partial p_i} + \frac{1}{4}Q(\widehat{p},\widehat{p})$. But $\frac{\partial^2 Q}{(e_i)^2} \leq 0$ and $\frac{\partial Q}{\partial e_i} < 0$, and therefore the RHS in (4) decreases with $\widehat{e}$. Finally, the fact that $Q(\widehat{e},\widehat{e})$ is 1 for $\widehat{e} = 0$ and 0 for $\widehat{e} = 1$ together with the fact that the agents' cost function satisfies the Inada conditions ensure that there exists a unique symmetric equilibrium.

(b) Each agent's marginal benefit of effort in the RHS of (4) can be written as $-Q'\left(R^i + R^j R^i \frac{1}{\eta}\right)$, where $R^i = \frac{e_i}{E}$ is agent i's proportional effort and $\eta_i = \frac{Q'}{Q}e_i$ is the elasticity of the product with respect to e_i. The incentive to exert effort thus depends on whether $R^i \frac{1}{\eta_i}$ is greater, equal to, or lower than 1. Now, in the socially-optimal outcome, agents exert equal efforts and therefore $R^i = \frac{1}{2}$. Whether agents exert socially-optimal efforts thus depends on whether η_i is greater, equal to, or smaller than $\frac{1}{2}$ or, equivalently, whether $\eta = \sum_{i=1,2}\eta_i$, the total elasticity of the product with respect to the agents' effort levels, is greater, equal to, or smaller than 1. In particular, if $\eta = 1$ (constant returns to scale), each agent's private marginal benefit of effort is equal to the social marginal benefit, Q'; agents thus have socially optimal incentives to exert effort. In contrast, if $\eta > 1$ (increasing returns to scale) agents exert too little effort in equilibrium and if $\eta < 1$ (decreasing returns to scale) agents exert too much effort. ■

Proposition 3 *The diligent's relative product share under a reciprocal allocation is lower than its proportional share if Q has decreasing returns to scale and non-decreasing differences and is greater than its proportional share if Q has increasing returns to scale and non-increasing differences.*

Proof. Suppose $e_d = \beta e_s$, where $\beta > 1$. Then the diligent's *proportional* share is $\frac{\beta}{1+\beta}$ and the slack's *proportional* share is $\frac{1}{1+\beta}$ (note that $\frac{\beta e_s}{\beta e_s + e_s} = \frac{\beta}{1+\beta}$). Under the RB rule, the slack's relative share is $\frac{1}{2}Q(e_s, e_s)/Q(e_s, e_d)$ and the diligent's relative share is the complement to 1. The slack's relative share under the RB rule is greater (lower) than its proportional share if $\frac{1}{2}Q(e_s, e_s)/Q(e_s, e_d) > (<)\frac{1}{1+\beta}$, or, equivalently, if $\frac{1+\beta}{2}Q(e_s, e_s) > (<)Q(e_s, e_d)$ (vice versa for the diligent). That is, the slack's relative share under the RB rule is greater (lower) than its proportional share if increasing one agent's effort from e_s to e_d (holding the other agent's effort fixed at e_s), increase the product by more (less) than $\frac{1+\beta}{2}$ times. We proceed by showing that, if Q has decreasing returns to scale and non-decreasing differences, the slack's relative share under the RB rule is greater than its proportional share.

Suppose Q has decreasing returns to scale and non-decreasing differences. Decreasing returns to scale imply that increasing both agents' efforts by $\beta - 1$ percent increases the expected product by less than $\beta - 1$ percent: $Q(e_d, e_d) < (\beta - 1)Q(e_s, e_s)$. Let $\overline{Q}(e_i, e_j) := \frac{1}{2}[Q(e_d, e_d) - Q(e_s, e_s)]$, where $i, j \in \{s, d\}$ and $i \neq j$, be the *average* increase in the product when both agents increase their efforts from e_s to e_d. Non-decreasing differences imply that $Q(e_i, e_j) - Q(e_s, e_s) < Q(e_d, e_d) - Q(e_i, e_j)$ and therefore that $Q(e_i, e_j) - Q(e_s, e_s) < \overline{Q}(e_i, e_j)$. This means that that, as one agent increases its effort from e_s to e_d, the corresponding increase in the product is lower than the half of the increase in the product when both agents increase their efforts from e_s to e_d. Together with the fact that Q has decreasing returns to scale, this implies that increasing one agent's effort by $\beta - 1$ percent increases the expected product by less than $(\beta - 1)/2$ percent. We thus have $\frac{1+\beta}{2}Q(e_s, e_s) > Q(e_s, \beta e_s)$. The proof of the case of Q with increasing returns to scale and non-decreasing differences follows naturally. ■

References

Aumann, Robert J. and Michael Maschler, Game Theoretic Analysis of a Bankruptcy Problem from the Talmud, *Journal of Economic Theory*, 36, 195-213 (1985).

Bonin, John P., Jones, Derek C., and Louis Putterman, AssociationTheoretical and Empirical Studies of Producer Cooperatives: Will Ever the Twain Meet? *Journal of Economic Literature*, 31(3), 1290-1320 (1993).

Bachrach, Yoram, Syrgkanis, Vasilis, and Milan Vojnovi´c, Efficiency and the Redistribution of Welfare, MSR-TR-2011-113 (2012)
(available at http://research.microsoft.com/pubs/155543/MSR-TR-2011-113.pdf).

Cornes, Richar and Roger Hartley, Joint Production Games with Mixed Sharing Rules (2004) (available at http://repec.org/esAUSM04/up.29210.1077802256.pdf).

Friedman Eric J., Strong monotonicity in surplus sharing. *Economic Theory,* 23, 643-658 (2004).

Israelsen, Dwight, Collectives, Communes, and Incentives, *Journal of Comparative Economics*, 4, 99-124 (1980).

Kornhauser, Lewis and Richard Revesz, "Sharing Damages among Multiple Tortfeasors," *Yale Law Journal*, 98, 831-884 (1989).

Kornhauser, Lewis and Richard Revesz, "Evaluating the Effects of Alternative Superfund Liability Rules," in Revesz, Richard L. and Stewart, R. (eds), Analyzing Superfund: Economics, Science, and Law (1995).

Leroux, Justin, "Cooperative Production under Diminishing Marginal Returns: Interpreting Fixed-Path Methods," *Social Choice and Welfare*, 29, 35–53 (2007).

Sen, Amartya K., Labour Allocation in a Cooperative Enterprise, *Review of Economic Studies*, 33, 361-371 (1966).

United States v. Reliable Transfer Co. 421 U.S. 397 (1975).