Rebels without a Cause
- Formal Model -*

Bernd Beber[†] Chris Blattman[‡]

July 2008

Abstract

Abduction as a means of recruitment is common among rebel groups in places such as northern Uganda, Sierra Leone, Nepal, and others. This paper asks what explains rebel groups' targeting of different population (and specifically age) groups for forcible recruitment. We present a formal model of the targeting choices made by rebel leaders, and show that they target individuals of intermediate ability because less able individuals contribute suboptimally to rebel activities while highly able individuals are difficult to retain. The extent to which leaders reward an abductee for effort contributed toward rebel objectives is increasing in ability, because they are constrained by having to make staying with the rebel unit at least as palatable as the (risky) outside option of an escape attempt. We test these hypotheses against a representative survey of 1016 households and 1360 male and female youth in northern Uganda, including 695 former abductees, which was implemented by Blattman and Annan (2007). We argue that age is a reasonable proxy for ability and show that the relationship between the likelihood of abduction, abduction length, and the likelihood of an abductee being a fighter (as opposed to serving in a non-combat support role) on the one hand and abductee age on the other hand can be characterized by an inverse U-shape. The extent to which abductees are subjected to threats or are forced to harm their family or friends is decreasing in age, and the degree to which abductees experience propaganda or are given rewards is increasing in abductee age.

*This paper is an incomplete version of "Rebels without a Cause: The Use of Coercion and Children in Guerrilla Warfare" providing some additional details about the formal model.

[†]Ph.D. candidate, Columbia University, bhb2102@columbia.edu

[‡]Assistant Professor, Yale University

1 The model

1.1 The structure of the game

Consider a game Γ with two actors, a rebel leader L and an abductee A. The abductee chooses how much effort a to contribute to rebel activities, while the rebel leader faces two choices: First, what is the optimal reward or punishment ρ given to the recruit, and second, what is the leader's optimal choice for productivity factor θ in recruitment? The rebel leader optimizes in the face of a principal-agent problem, i.e. she observes productive output but is unable to directly observe either the abductee's effort level or the random component in production, while the abductee is fully informed.[1]

Utilities are given by

$$u^A = r(\rho, \pi(a, \theta, \varepsilon)) - c^A(a) \quad \text{and}$$
$$u^L = \pi(a, \theta, \varepsilon) - c^L(\rho, \pi(a, \theta, \varepsilon)),$$

where output is given by production function π with non-negative function inputs a and θ and random component ε, the abductee derives utility from output priced at $\rho \in \mathbb{R}$ and disutility given by some cost function c^A of effort, and the rebel leader's utility is given by net output.[2] For notational convenience all indices are denoted as superscripts. Subscripts denote partial derivatives. In the following function inputs are usually not explicitly specified.

We assume that cost and production functions are positive-valued. All functions are continuous and thrice differentiable over their domain, production function π has positive first-order partials, and reward function r is increasing and concave in π. We also assume that $c^A_{aa} \geq 0$, i.e. the abductee's cost of providing additional effort is non-decreasing in the amount of effort already contributed. The recruit experiences diminishing returns to effort, i.e. $\pi_{aa} < 0$ and hence $r_{aa} < 0$.[3] We assume that $u^A_a > 0$ at $a = 0$, i.e. a recruit who is not showing any effort at all can expect to gain net utility if he increases his effort.

Both a non-zero amount of effort and ability are required in order for the leader

[1] We use female pronouns to refer to the leader and male pronouns to refer to the abductee.

[2] As is common in the literature on principal-agent models, we assume that utility is additively separable into revenue and costs (Bolton and Dewatripont, 2005: 130,142).

[3] This is an assumption on production function π, since $\pi_{aa} < 0$ implies $\frac{\pi_{aa}}{\pi_a \pi_a} < \frac{-r_{\pi\pi}}{r_\pi}$ and so $r_{aa} = r_{\pi\pi}\pi_a\pi_a + r_\pi \pi_{aa} < 0$.

to accrue utility, that is $\pi = 0$ at $a = 0$ or $\theta = 0$. In order for the abductee to derive utility from an increase in production, some rewards or punishments have to be offered, i.e. $r_\pi = 0$ if $\rho = 0$. We assume that "market entry" for the leader is optimal, i.e. $c^L_\pi < 1$ at $\pi = 0$.

The rate at which the leader's cost rises with output is increasing in the incentives given to the abductee, i.e. $c^L_{\pi\rho} > 0$, and the rate at which the abductee's benefits from rebel activities rise with effort is increasing in his productivity, i.e. $r_{a\theta} \geq 0$.[4] For the leader's cost function, we have $c^L_\rho < (>) 0$ for $\rho < (>) 0$, and $c^L_\rho = 0$ for $\rho = 0$, which means the leader's costs are minimized when she neither rewards nor punishes.

We require that a participation constraint is met for the abductee, which is that the recruit's utility exceeds some reservation point $\underline{u}(\theta)$, where $\underline{u}$ is a function of his productivity level θ. We assume that if $r - c^A \leq \underline{u}$ at some $\tilde{\theta}$, then $r - c^A < \underline{u}$ for all $\theta > \tilde{\theta}$, which ensures that if the participation constraint is not met for some level of productivity, then it won't be met for higher productivity levels either. Finally, we assume $r - c^A > \underline{u}$ at $\theta = 0$ and $r - c^A < \underline{u}$ as $\theta \to \infty$, i.e. the participation constraint is met for abductees of the lowest type, but not for abductees of the highest type. The former rules out the case in which the participation constraint is never met, and the latter dispenses with the case in which no participation constraint has to be met. (Another way of stating the same assumption is to require $r - c^A > \underline{u}$ at $\theta = 0$ as well as $\bar{\underline{u}}_\theta > \bar{r}_\theta$, where the upper bar denotes the average over the domain of θ, i.e. an increase in productivity will in expectation boost the abductee's reservation price more than it will increase his benefits accrued from rebel service.)

1.2 Related literature

The model presented here draws on contract theory and is technically reminiscent of principal-agent models of adverse selection (where the principal does not observe agent type prior to contracting) and moral hazard (where agent actions after contracting are unobservable to the principal).[5] With respect to the former, our model recalls in particular the literature on countervailing incentives. In one strand of

[4]More specifically, what this requires is that $r_{a\theta} = r_{\pi\pi}\pi_\theta\pi_a + r_\pi\pi_{a\theta} \geq 0$, which we can write as $\frac{\pi_{a\theta}}{\pi_\theta\pi_a} \geq \frac{-r_{\pi\pi}}{r_\pi}$. We assume that $\pi_{a\theta}$ is positive and sufficiently large for this condition to hold, i.e. effort and ability are sufficient complements. In other words, we write assumption $r_{a\theta} \geq 0$ in terms of reward function r for convenience, but are actually restricting production function π.

[5]For a non-technical overview of the application of principal-agent models in political science, see Miller (2005).

this literature, type-dependent participation constraints ensure that agents across the type space do not all have an incentive to overstate (or understate) their type, where an agent's type in this literature is typically his marginal cost of production (Lewis and Sappington, 1989; Maggi and Rodríguez-Clare, 1995; Jullien, 2000): An agent's incentive to overstate costs to boost revenue can be balanced by a reservation value that is decreasing in costs.[6] Like the models in that literature, we propose that the agent's reservation value is endogenous to his type. The principal in our model, however, is more high-powered: First, the leader observes agent type, which removes the principal's key challenge of adverse selection and permits her to select the first-best schedule of reward or punishment for each agent. Second, the leader makes a choice over type, which means she can select her most preferred among the first-best allocations. The question then is not what kind of incentives the agent faces in revealing his type (since his type is known), but rather what kind of countervailing incentives the endogenous participation constraint generates for the principal in type selection.[7]

Since the leader sets agent type, our endogenous participation constraint binds with equality for all types, in contrast to many adverse selection models. Typically, the participation (or individual rationality) constraint binds only at the low end of the type distribution (see for example Boone and Bovenberg, 2004).[8] But since the leader in our model can discriminate among abductees of differing ability (and given the natural assumption that ability and effort are sufficiently complementary rather than substitutable in production, that is $\pi_{a\theta} >> 0$),[9] the rebel leader will calibrate rewards and punishment to just meet an abductee's reservation utility. For the same reason, we see full participation of (selected) agents in our model. Any adjustments of labor supply occur at the intensive (effort level) as opposed to extensive margin (recruitment level).[10]

[6]This assumes that the agent's reservation utility is concave in marginal cost, as Maggi and Rodríguez-Clare (1995) have shown.

[7]Another difference to most models in this literature (e.g. Jullien, 2000) is that the principal's utility in our model is affected by agent type, which we argue is appropriate in this substantive context. In a typical adverse selection model, this could create the problem that while utility inputs have to be verifiable, the agent's type is not. In our model, this problem does not arise because the principal observes agent type.

[8]For an exception, see Rochet and Stole (2002), where agents vary along two dimensions (search costs and skills) and a binding participation constraint can be derived for each type.

[9]For a historically grounded discussion of complementarity and substitutability of inputs, see Weber (2005).

[10]For an example of adjustment at the extensive margin, see Boone and Bovenberg (2004), where search costs and a participation constraint lead to low-skill workers dropping out of workforce.

The (in this context misleadingly named) moral hazard problem impinges more seriously on the leader's optimization.[11] It prevents the leader from obtaining her first-best outcome for two reasons. First, agent effort is not observable to the leader. Second, the agent's utility is sensitive to risk, which means that the leader can't "sell" productive output to the agent upfront and keep a risk-free, constant share for herself. Due to the agent's risk aversion, he will not optimally be designated the output's residual claimant (Laffont and Martimort, 2002: 153).

Note, though, that even if the abductee had risk-neutral preferences, this would not imply that only he is exposed to risk, as is the case for the first-best in other principal-agent models.[12] In our model the leader and abductee share risk by construction, because the leader's wealth is just equal to (stochastic) output—we do not allow the leader to have the abductee operate as a "franchise", with the leader receiving a constant fee independent of productive output. This is a natural assumption given our substantive application, since rebel leaders cannot generally insulate themselves entirely from the risks that come with a battlefield confrontation. Running a rebellion is not risk-free.

We now turn to the wage optimization on the part of the principal. The leader's problem is to set a wage for each possible realization of productive output, i.e. π. Since π is continuous, the leader has to optimally choose a wage function. Optimization over a function space is complex, however,[13] and we make the problem tractable by limiting the leader's optimization to a specific (and known) array of curves, which is defined over an array parameter ρ. Instead of allowing the leader to choose any function $r(.)$, we fix $r(.)$ and have the leader optimize over scalar ρ. This makes our analysis significantly less general, although it is still more general than a game with a binary outcome space. The main disadvantage is that we impose attributes on the wage function that may not be shared by the function that the principal would choose optimally. Specifically we require that r is increasing and concave in π. The optimal wage function does in fact have this property in models with a binary outcome, but this is not generally true with a continuous outcome. While it has been shown that wages cannot be uniformly decreasing in productive output, they may be locally decreasing (Grossman and Hart, 1983, cited in Salanié, 2005: 129).

[11] See Ferejohn (1986) for a classic treatment of the moral hazard problem in political science.
[12] See Shavell (1979: 59) for the basic argument.
[13] For examples, see Page (1987), cited in Salanié (2005: 139), and Shavell (1979).

We should mention two applications of the principal-agent setup that are particularly relevant to our model. First, Scott Gates' (2002) model of rebel recruitment addresses the puzzle why forcibly recruited soldiers would contribute to rebel activities and shows, similar to us, how recruits maneuver to reap rewards and avoid punishment. There are three key differences to our approach, however. First, Gates argues that abducted recruits have a reservation value that is equivalent to death (which in his model is normalized to a constant, 0). But while it is true that a LRA abductee faced a real possibility of being killed if he tried to escape, death was generally a probabilistic component of his outside option—it was not a certainty. Gates' assumption means that abductees are more pliant in his model than they are in ours. Second, we add verisimilitude to our model by allowing reservation values to be endogenous to agent type, which implies that the extent to which the leader has to treat an abductee favorably is increasing in the abductee's skill level. Third, Gates focuses on the contract offered to recruits, but his model does not speak to which kinds of individuals are most likely to be targeted for recruitment.

The second application is Chwe's (1990) paper on pain in a principal-agent setup. His main result is that a sufficient increase in a worker's reservation value means that the optimal contract will provide money rather than pain. Children, who have poor outside options, are then particularly likely to suffer punishment. Chwe shows that the principal would never use both money and pain after a given outcome, which we incorporate as part of the structure of the game: Here the leader provides a real-valued, unidimensional incentive, which can be negative or positive but obviously never both. As with Gates' model, the main advantage our approach offers to Chwe's is the fact that we treat reservation values as endogenous to agent type. Chwe's analysis emphasizes the role of the participation constraint in explaining why some workers were whipped, but takes the worker's reservation utility as fixed. We generalize this feature of the principal-agent interaction.

1.3 A general result

We want to show that only an interior solution exists for θ, i.e. setting θ equal to its largest or smallest possible value is not optimal for the leader. First, the abductee's maximization problem is implicitly solved by choosing a such that $r_a = c_a^A$. For incentive compatibility, this constraint enters the leader's optimization problem.[14]

[14] In principal-agent models, incentive compatibility usually refers to the constraint under which agents reveal private information truthfully. Here we follow Chwe (1990) and use the term incentive

We also require that the participation constraint is met. The leader's non-linear programming problem then is to find θ and ρ that maximize u^L such that

$$r_a = c_a^A \quad \text{and} \tag{1}$$

$$r - c^A \geq \underline{u}. \tag{2}$$

First, we show that the abductee can solve equation (1) for a, given θ and ρ. Recall that $c_{aa}^A \geq 0$ and $r_{aa} < 0$, and hence as $a \to \infty$, $u_a^A < 0$. We also assumed that $u_a^A > 0$ at $a = 0$. But then by the Intermediate Value Theorem a solution a^* must exist for $u_a^A = 0$, i.e. equation (1).

Since we can find some a^* that will ensure incentive compatibility for any θ and ρ chosen by the rebel leader, we can write a^* as a function of θ and ρ and in turn drop constraint (1) from the leader's optimization problem. The leader's problem now reads

$$\max_{\theta,\rho} u^L(\theta, \rho, a^*(\theta, \rho)) \quad \text{s.t. } r - c^A \geq \underline{u},$$

with Lagrangian

$$\mathscr{L} = \pi - c^L + \lambda(r - \underline{u} - c^A)$$

and first-order conditions

$$\mathscr{L}_\rho + \mathscr{L}_a a_\rho^* = 0, \text{ and}$$
$$\mathscr{L}_\theta + \mathscr{L}_a a_\theta^* = 0.$$

Since a^* is implicitly defined by $u_a^A = 0$, where u_a^A is continuous around solution (a^*, θ^*, ρ^*) (because u^A is twice differentiable in the three parameters over their domains) and $u_{aa}^A \neq 0$ (because $c_{aa}^A \geq 0$ and $r_{aa} < 0$), we can use the Implicit Function Theorem to compute a_ρ^* and a_θ^*. Taking derivatives, the first-order conditions become

$$-c_\rho^L + \lambda r_\rho + (\pi_a - c_a^L)\left(\frac{-r_{a\rho}}{r_{aa} - c_{aa}^A}\right) = 0, \text{ and} \tag{3}$$

$$\pi_\theta - c_\theta^L + \lambda(r_\theta - \underline{u}_\theta) + (\pi_a - c_a^L)\left(\frac{-r_{a\theta}}{r_{aa} - c_{aa}^A}\right) = 0. \tag{4}$$

compatibility to refer to the solution of the agent's optimization problem.

where the Lagrange multiplier $\lambda \geq 0$.[15] We now show that the participation constraint binds at any solution.

The intuition behind the proof is as follows. Suppose the participation constraint is not active. That means that the leader can choose the abductee's skill level without having to worry about picking someone so capable he would leave, and she will set the agent's skill level (which enters the leader's utility through production) so as to equalize marginal revenue and marginal cost and realize an optimal level of production. Now the leader still has to optimize the incentive schedule offered to the agent and, by way of these incentives, the agent's effort level. So what effort level will the leader aim for? Note that effort only enters the leader's utility through production—but production is already optimal by way of the skill level that the leader selected! This means that any effort level will satisfy the leader, who in turn minimizes costs by offering neither rewards nor punishment. But without incentives, the abductee provides no effort and no output is generated, which is not optimal for the leader by assumption and hence a contradiction.

Lemma 1. *At any solution of the leader L's optimization problem, the participation constraint binds with equality, i.e. $r - c^A = \underline{u}$.*

Proof. Suppose to the contrary that there is some solution at which the participation constraint does not bind. Then $\lambda = 0$, and since $c^L_a = c^L_\pi \pi_a$ as well as $c^L_\theta = c^L_\pi \pi_\theta$, we can factor equation (4) and have

$$(1 - c^L_\pi)\left(\pi_\theta + \pi_a\left(\frac{-r_{a\theta}}{r_{aa} - c^A_{aa}}\right)\right) = 0. \tag{5}$$

Since $r_{a\theta} \geq 0$ by assumption, the second factor is positive and $c^L_\pi = 1$ at any solution. Equation (3) in turn reduces to $-c^L_\rho = 0$, which implies $\rho = 0$. But $r_\pi = 0$ at $\rho = 0$, which means the abductee has no incentive to exert effort and sets $a = 0$. Hence $\pi = 0$, which implies $c^L_\pi < 1$ (because market entry is optimal), a contradiction. □

This result is not as straightforward as it may appear, because of the leader's ability to punish. Suppose some abductee will stay with the leader even if subjected to punishment. In equilibrium the leader will then induce effort by punishing the abductee as much as she can without the abductee trying to escape. This is the case even though the leader could provide rewards *at exactly the same cost.* The

[15] $\mathscr{L}_a$ is equal to $\pi_a - c^L_a + \lambda r_a - \lambda c^A_a$, which simplifies to $\pi_a - c^L_a$ because of incentive compatibility.

intuition is that if the leader is providing rewards, then she will abduct someone skilled who requires rewards to stay.

Since the participation constraint binds at any solution, it follows immediately that the extent to which the abductee's equilibrium payoff varies with his type can be characterized by how his reservation value depends on skill level. If reservation utility is a monotonically increasing function of productivity, then the abductee's equilibrium payoff is monotonically increasing in skill. This implies in particular that types above some skill threshold do not experience punishment.

Corollary 1. *Define $u^A \leq \hat{u}^A$ as punishment payoff and $u^A > \hat{u}^A$ as reward payoff. If $\underline{u}(\theta)$ is monotonically increasing in θ, then there is some threshold $\hat{\theta}$ such that agents of type $\theta \leq \hat{\theta}$ are punished and agents of type $\theta > \hat{\theta}$ are rewarded in equilibrium.*

Proof. If $\underline{u}_\theta > 0$, then by Lemma 1 u^{A*} is also monotonically increasing in θ and hence there exists a unique $\hat{\theta}$ such that the claim holds. □

We now turn to the optimal skill level selected by the leader. We have three equations with three unknowns: First-order conditions (3) and (4) and the participation constraint. Our strategy is to show that we can solve each equation separately for a different parameter and then use a fix point theorem to establish that a solution exists for the system of equations as well.

Proposition 1. *In game Γ, a solution of the leader L's optimization problem for productivity parameter θ exists and is located in the interior of the domain of θ.*

Proof. The leader's optimization problem is given by first-order conditions (3) and (4) and participation constraint $r - c^A = \underline{u}$. We first solve equation (4) for λ, which yields

$$\lambda = \frac{(1 - c^L_\pi)\left(\pi_\theta + \pi_a\left(\frac{-r_{a\theta}}{r_{aa} - c^A_{aa}}\right)\right)}{\underline{u}_\theta - r_\theta}. \tag{6}$$

By assumption, $\underline{u}$ intersects only once with $r - c^A$, from below, and hence $\underline{u}_\theta > r_\theta$ at any solution. In turn $\lambda \geq 0$, as required, for any $c^L_\pi \leq 1$.

We solve equation (3) for ρ. Note that the LHS of (3) is continuous in ρ. For some sufficiently large $|\rho|$, where ρ can be either positive or negative, we have $c^L_\pi = 1$ (i.e. units of output can be priced so aggressively that an increase in production is not worthwhile for the leader). This means $\lambda = 0$ from (6) and $\pi_a - c^L_a = \pi_a(1 - c^L_\pi) = 0$ and so equation (3) simplifies to $-c^L_\rho$, which is negative for $\rho > 0$ and positive for $\rho < 0$. But then some ρ^* that solves (3) must exist, by the Intermediate Value Theorem.

Third, we solve constraint (2) for θ. From Lemma 1, we know that $r - c^A = \underline{u}$ at any solution, where both sides of the equation are continuous in θ. Since $r - c^A > \underline{u}$ at $\theta = 0$ and $r - c^A < \underline{u}$ as $\theta \to \infty$, an interior solution θ^* must exist, again by the Intermediate Value Theorem.

We can now define the functions $f^\lambda(\rho^*, \theta^*)$, $f^\rho(\lambda^*, \theta^*)$, and $f^\theta(\lambda^*, \rho^*)$, where f^i gives the set of optimal values for i, given optimal values for parameters $j \neq i$. Let f be the Cartesian product of these functions. Function f is a mapping from $S \to S$, where S is the Cartesian product of the support of λ, ρ, and θ. Since $\lambda, \theta \in \mathbb{R}^+$, and $\rho \in \mathbb{R}$, we know that S is non-empty, compact, and convex (specifically, a convex subset of $\mathbb{R}^3$). The first three steps of the proof showed that $f(s)$ is non-empty for all $s \in S$. Since u^A and u^L are thrice differentiable in their inputs, functions f^i are continuous, and in turn function f is continuous. Then all conditions for Brouwer's fixed point theorem are met, and a fixed point exists, i.e. a solution for the leader's optimization problem exists.

Finally, we verify that a corner solution does not exist, which follows from Lemma 1 and the fact that $r - c^A \neq \underline{u}$ at either $\theta = 0$ or as $\theta \to \infty$. If $\theta = 0$, the abductee's reservation values is arbitrarily small and the participation constraint does not bind, but then the leader has an incentive to move her productivity target upward. Similarly, for sufficiently large θ, the abductee's reservation price is larger than the payoff he can receive from the leader, which means it is not a feasible solution. □

The key insight of Proposition 1 is that the leader arrives at an interior solution for θ^* because of a participation constraint, without any assumptions being imposed as to the concavity of u^L over θ.

1.4 An example with explicit functions

In order to develop an intuition for our main result, consider the following example with explicit functions. Suppose we have a simple production function $\pi = \theta a \varepsilon$ with single input a and "technology" parameter θ.[16] Assume that random variable ε is distributed Gamma with mean 1 and standard deviation s (i.e. with shape $1/s^2$ and scale s^2).[17] Utilities are given by

$$u^A = \begin{cases} \rho e^{-\phi\pi} - a^2 & \text{for } \rho \leq 0 \\ \rho(1 - e^{-\phi\pi}) - a^2 & \text{for } \rho > 0 \end{cases} \quad \text{and}$$

$$u^L = \pi - \rho^2 \pi,$$

where ϕ is the abductee's coefficient of absolute risk aversion.[18] For the participation constraint we will assume

$$\underline{u}(\theta) = \frac{(\theta - 1)^3}{\theta}.$$

Figures (1)–(3) illustrate some convenient properties of these utility functions, which reflect the assumptions of the general model.[19] We assume $\phi = 1$ for this set of figures. First, as illustrated in Figure (1), the participation constraint is satisfied for the lowest-productivity individual, but not for the abductee with the highest productivity.

The abductee's utility is strictly increasing in productive output, as shown in figure (2). The sign of u^A, that is whether the abductee is rewarded or is punished, is determined by the sign of the incentive scheme ρ, i.e. it depends on whether the rebel leader puts a rewards or a punishment program in place. The extent to which the abductee's utility is increasing in productive output is a function of the value of the reward, or equivalently, the severity of punishment.

Figure (3) shows how the abductee's utility changes with shifts in the leader's payout scheme, for different levels of productive output π (or different levels of effort, if we fix θ and ε). The graph illustrates that both an increase in output and a move away from punishment toward rewards leaves the abductee better off.

[16]Parameter θ represents production technology. In the language of principal-agent games, technology is captured by function π, which relates actions unobservable to the principal to verifiable output.

[17]Gamma distributed performance ensures that production is positive-valued, in contrast to a normal distribution.

[18]We define the agent's risk aversion vis-à-vis stochastically determined output π, so $\phi = -\frac{u^A_{\pi\pi}}{u^A_{\pi}}$.

[19]Utility function u^A is not in fact differentiable at $\rho = 0$, but this could be accomplished by modifying it close to 0 (Chwe: 1990, 7).

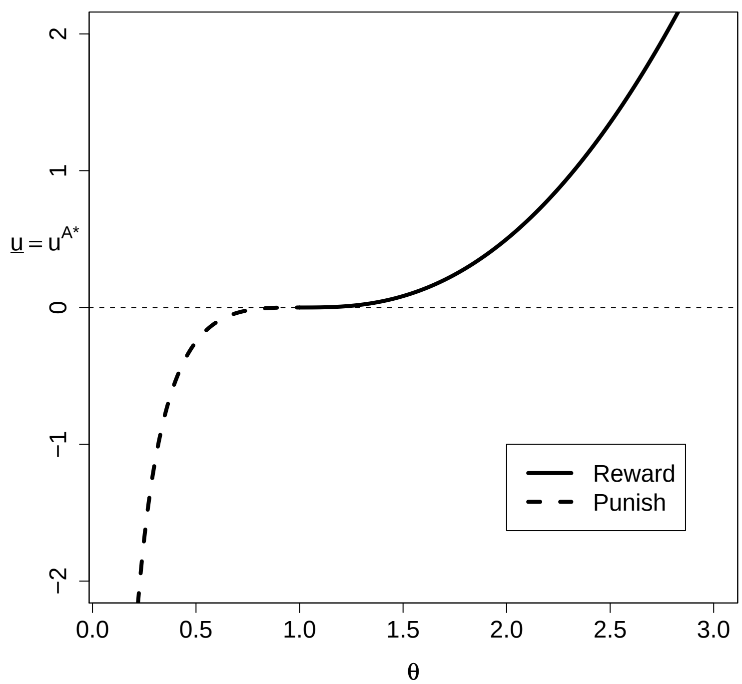

Figure 1: Abductee's reservation utility as a function of his productivity

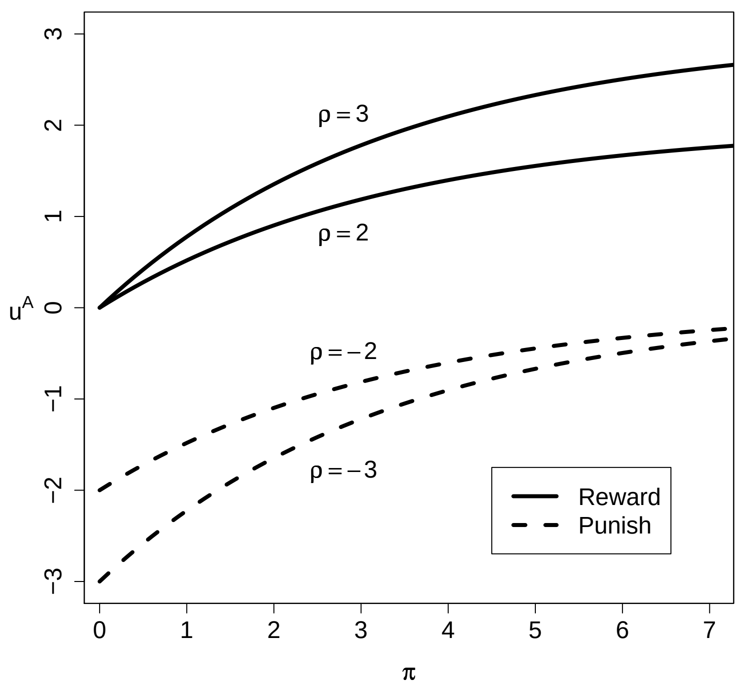

Figure 2: Abductee's utility as a function of productive output

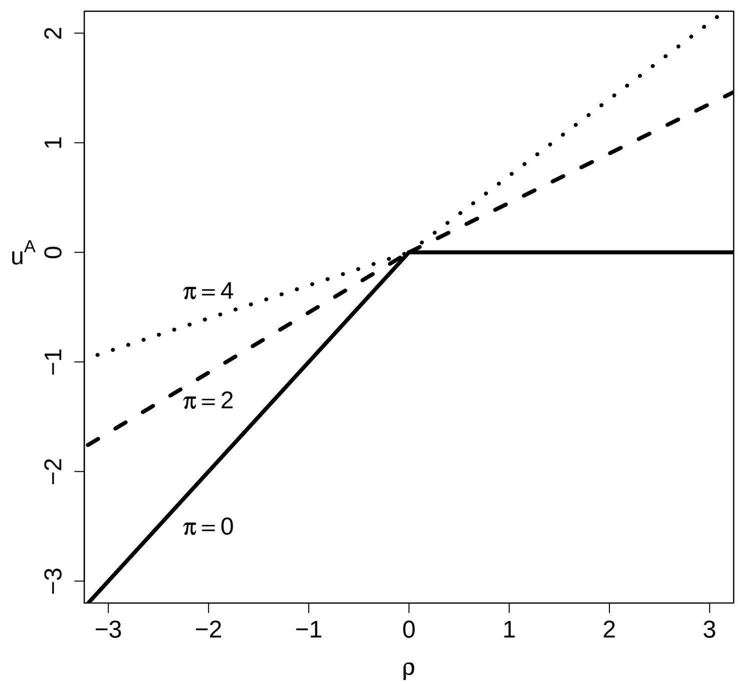

Figure 3: Abductee's utility as a function of the reward/punishment schedule

1.4.1 Payoffs when the leader punishes

We solve this example in two parts. First, let $\rho \leq 0$, i.e. the leader punishes the abductee. Since the leader does not observe stochastic component ε, she considers the agent's expected utility $\mathrm{E}(u^A) = \rho \mathrm{E}(e^{-\phi\pi}) - a^2$. We have

$$\begin{aligned}\mathrm{E}(e^{-\phi\pi}) &= \int_0^\infty e^{-\phi\pi} f(\varepsilon)\, \mathrm{d}\varepsilon \\ &= \frac{1}{(1+\phi\theta a s^2)^{1/s^2}},\end{aligned}$$

where $f(.)$ is the probability density function of the Gamma distribution.

Then incentive compatibility requires $-\phi\theta\rho(1+\phi\theta a s^2)^{-s^{-2}-1} = 2a$ and hence

$$\rho = \frac{-2a}{\phi\theta}(1+\phi\theta a s^2)^{1+1/s^2}. \tag{7}$$

The participation constraint gives us $\frac{\rho}{(1+\phi\theta a s^2)^{1/s^2}} - a^2 = \frac{(\theta-1)^3}{\theta}$.[20] We can then substitute using equation (7) and eliminate ρ, which leaves us with

$$a = \frac{-\theta + \sqrt{\theta^2 - \phi^2(1+2s^2)\underline{u}}}{\phi(1+2s^2)}. \tag{8}$$

Using equations (7) and (8) we can write u^L in terms of θ and maximize in one variable, over a grid of values for s and ϕ. For example, with $s = .5$ and $\phi = 3$, we find $\theta^{**} \approx .55$, $a^{**} \approx .23$, $\rho^{**} \approx -.43$, $u^{L**} \approx .10$, and $u^{A**} \approx -.16$.

[20] We use the result from Lemma 1, which requires $r_{a\theta} \geq 0$. This assumption is valid only for $\pi \leq 1$ in this example. However, a modified version of the proof still holds: By taking derivatives of the example's explicit functions and substituting in equation (5), we have

$$(1 - c^L_\pi)\left(a + \theta\left(\frac{(\pi-1)|\rho|e^{-\pi}}{-\theta^2|\rho|e^{-\pi} - 2}\right)\right) = 0.$$

Setting the second factor equal to 0 and solving for ρ yields $|\rho| = \frac{-2a}{\theta e^{-\pi}}$, a contradiction. Hence at any solution of equation (5), we have $c^L_\pi = 1$, as required for the rest of the proof. As we show below, $\pi^* < 1$ in any case.

1.4.2 The reward equilibrium

Now consider the case where $\rho > 0$. From incentive compatibility we have

$$\rho = \frac{2a}{\phi\theta}(1 + \phi\theta a s^2)^{1+1/s^2}.$$

We can leverage the participation constraint to implicitly define a in terms of θ, which yields

$$\frac{2a}{\phi}\left[(1 + \phi\theta a s^2)^{1+1/s^2} - (1 + \phi\theta a s^2)\right] - a^2\theta - (\theta - 1)^3 = 0.$$

Unlike in the previous case, we cannot solve algebraically for a and then substitute into the objective function u^L, but we can compute u^L numerically and maximize over θ.[21] With $s = .5$ and $\phi = 3$, for example, we locate $\theta^* \approx 1.79$, and hence $a^* \approx .27$, $\rho^* \approx .48$, $u^{L*} \approx .37$, and $u^{A*} \approx .27$.

Since $u^{L*} > u^{L**}$, in this case the leader strictly prefers to reward rather than punish an optimally recruited individual. Since the leader moves first by way of setting the incentive scheme and choosing a recruitment strategy, the unique equilibrium of the game is given by θ^*, a^*, and ρ^*. The leader here abducts relatively productive individuals who have to be treated relatively well (or else they would desert), but the key point is that even in this case the leader does not choose to target the most productive individuals that could be abducted.

Figure 4 shows the rebel leader's payoff from targeting individuals of different productivity levels. Clearly, choosing $\theta^* \approx 1.79$ yields the highest utility for $s = .5$ and $\phi = 3$.

Figure (5) illustrates the equilibrium effort level exhibited by abductees of differing productivity, again for $s = .5$ and $\phi = 3$. Effort is lowest when the rebel leader offers neither rewards nor punishment, and it is increasing in either direction: As productivity increases, it becomes more profitable for the leader to reward input, and in turn effort goes up. As productivity falls below $\theta = 1$, on the other hand, the abductee's reservation value drops and punishment becomes an increasingly feasible instrument to enforce greater effort.

[21] `R` code is available from the authors.

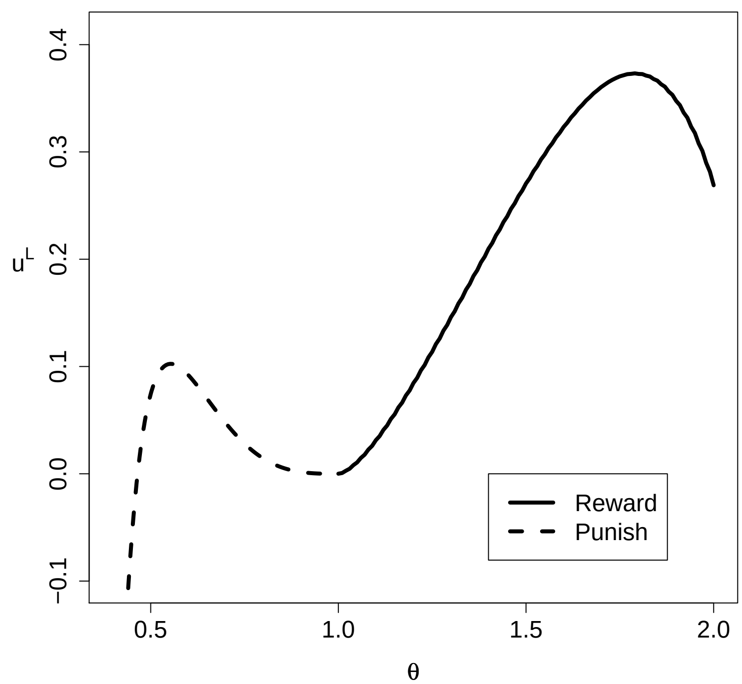

Figure 4: Leader's utility in terms of abductee productivity

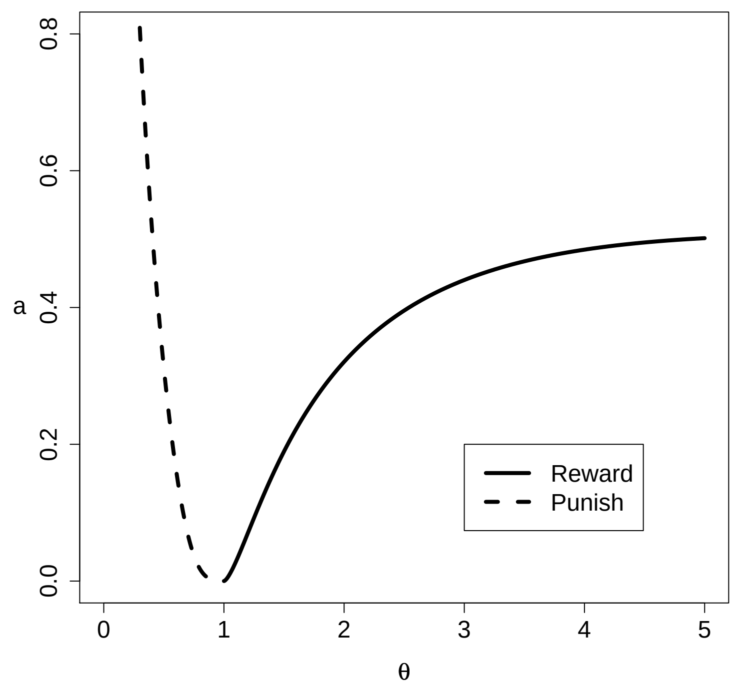

Figure 5: Abductee's effort in terms of his productivity

1.4.3 When is punishment an equilibrium strategy?

Figure 4 shows that the leader provides rewards in equilibrium, conditional on a high level of risk aversion $\phi = 3$ and a non-trivial degree of uncertainty $s = .5$. In fact the leader always provides rewards given the utilities we have considered so far in this example, regardless of the specific values we choose for ϕ and s. The reason is that so far we have only considered the situation in which punishment is as expensive to the leader as the provision of rewards. We now generalize this feature of the model.

Let $u^L = \pi - k\rho^2\pi$, where $k \in \mathbb{R}$ and $k = 1$ for $\rho > 0$. If $k < 1$ for $\rho < 0$, the extent to which the leader's cost function is increasing in rewards is greater than (or equal to) the extent to which it is increasing in punishment. If $k = 0$, it does not cost the leader anything to punish (and if $k < 0$, the leader derives pleasure from it). If $k = 1$, inflicting punishment is as costly as the provision of rewards. We calculate equilibrium values as before, and find that punishment can indeed be an equilibrium strategy for sufficiently small k, provided both risk aversion (measured in terms of ϕ) and uncertainty (in terms of s) are low. Figure 6 illustrates this point. In this example, $s = .5$, $\phi = .3$, and $k = .001$.

The model describes three ways in which a rewards-driven equilibrium could be obtained. First, if abductees are sufficiently averse to the risks associated with participating in rebel activities, the leader will target relatively productive individuals and provide rewards. The same will obtain if, second, production uncertainty is large. If the leader is sufficiently unsure about what the outcome of her plan will be, she will provide rewards rather than inflict punishment. Finally, the leader's optimal strategy won't involve punishment if a punishment regime is costly to maintain.

1.5 Discussion and implications

Three additional comments related to figure (4) are in order. First, although the utility function depicted in (4) is bimodal, this implies a unimodal empirical distribution. In our example, we expect the rebel leader to choose θ^*; we would not expect θ^{**} to play a role in corresponding observational data. In turn we should observe a single mode at θ^*.[22]

In the general case, the empirical distribution that we would expect around θ^* does not follow from the graph of u^L, but rather from production error ε. If we

[22]There is a special, but non-generic case in which $\theta^* = \theta^{**}$.

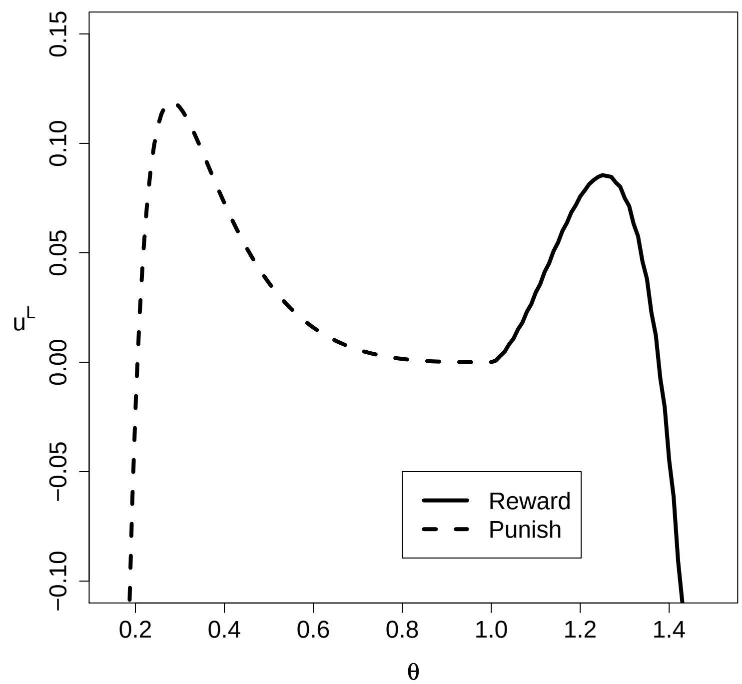

Figure 6: Leader's utility when punishment is cheap

assume that the rebel leader's prior belief about ε is drawn from some symmetric distribution, then the rebel leader's optimal productivity target will be drawn from a symmetric distribution around θ^*. In fact the only role that ε plays in our model is to introduce some uncertainty, by way of prior belief distributions, over equilibrium parameter values. The abductee, although fully informed about production error ε, never sends a signal of consequence to the rebel leader, because the incentive schedule is fully determined prior to any action by the abductee.

Second, figure (4) helps to illustrate the effect of imposing an additional budget constraint on the rebel leader. Recall that both punishments and rewards are costly for the rebel leader, and an active budget constraint will force the leader to move closer to not taking either action (which corresponds to the point at which the solid and the dashed lines meet in figure (4)). In turn, if the leader is punishing in equilibrium, this means that the optimal target productivity will shift upward, but if the leader is providing rewards in equilibrium, then her productivity target will shift downward. If age is a useful proxy of productivity,[23] then a punishment-oriented rebel leader would respond to budget pressures by targeting fewer children; a rewards-oriented leader would respond by targeting more children.

Finally, consider the effect of a change in the mapping from θ to $\underline{u}$. Suppose, for instance,

$$\underline{u}(\theta) = \frac{(\theta - \theta_0)^3}{\theta},$$

where $\theta_0 = 1$ in our example. Reservation prices will increase the more steeply in productivity, the smaller θ_0. If age is a proxy for θ, then we could think of a society's economic or educational opportunities as a proxy for θ_0. From figure (1), we can see immediately that any decrease in θ_0 will leave abductees better off: Their utility (if abducted) will be higher than it would be otherwise, because rebel leaders now have to treat them better in order to retain them. But at the same time, exactly because high-productivity abductees have now become more difficult to retain, rebel leaders will take to conscripting more child soldiers.

[23]See Skirbekk (2004) for on overview of the literature on the relationship between productivity and age. Skirbekk shows that the average adult's productivity is increasing in age until it peaks when he or she is between thirty and fifty years old, depending on occupation. Since about 93% of the survey respondents we analyze in the empirical section of this paper are 30 years or younger, we find it reasonable to take age as an indicator of increasing productivity.

References

Christopher Blattman and Jeannie Annan. The consequences of child soldiering. Households in Conflict Network working paper 22, 2007.

Patrick Bolton and Mathias Dewatripont. *Contract Theory.* Cambridge, MA: The MIT Press, 2005.

Jan Boone and Lans Bovenberg. The optimal taxation of unskilled labor with job search and social assistance. *Journal of Public Economics*, 88(11):2227–2258, 2004.

Michael Suk-Young Chwe. Why were workers whipped? Pain in a principal-agent model. *The Economic Journal*, 100(403):1109–1121, 1990.

John E. Ferejohn. Incumbent performance and electoral control. *Public Choice*, 50 (1–3):5–25, 1986.

Scott Gates. Recruitment and allegiance: The microfoundations of rebellion. *Journal of Conflict Resolution*, 46(1):111–130, 2002.

Sanford J. Grossman and Oliver D. Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1):7–45, 1983.

Bruno Jullien. Participation constraints in adverse selection models. *Journal of Economic Theory*, 93(1):1–47, 2000.

Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model.* Princeton, NJ: Princeton University Press, 2002.

Tracy R. Lewis and David E. M. Sappington. Countervailing incentives in agency problems. *Journal of Economic Theory*, 49(2):294–313, 1989.

Giovanni Maggi and Andrés Rodríguez-Clare. On countervailing incentives. *Journal of Economic Theory*, 66(1):238–263, 1995.

Gary J. Miller. The political evolution of principal-agent models. *Annual Review of Political Science*, 8:203–225, 2005.

Frank H. Page, Jr. The existence of optimal contracts in the principal-agent model. *Journal of Mathematical Economics*, 16(2):157–167, 1987.

Jean-Charles Rochet and Lars A. Stole. Nonlinear pricing with random participation. *The Review of Economic Studies*, 69(1):277–311, 2002.

Bernard Salanié. *The Economics of Contracts: A Primer*. Cambridge, MA: The MIT Press, second edition, 2005.

Steven Shavell. Risk sharing and incentives in the principal and agent relationship. *The Bell Journal of Economics*, 10(1):55–73, 1979.

Vegard Skirbekk. Age and individual productivity: A literature survey. In Gustav Feichtinger, editor, *Vienna Yearbook of Population Research 2004*, pages 133–153. Vienna: Austrian Academy of Sciences Press, 2004.

Christian E. Weber. Edgeworth on complementarity, or Edgeworth, Auspitz-Lieben, and Pareto de-homogenized. *History of Political Economy*, 37(2):293–307, 2005.