

Nash Equilibria for Weakest Target Security Games with Heterogeneous Agents^{*}

Benjamin Johnson^{a,d}, Jens Grossklags^b, Nicolas Christin^{a,c}, and John Chuang^d

^aCyLab, Carnegie Mellon University

^bCollege of Information Sciences and Technology, Pennsylvania State University

^cInformation Networking Institute, Carnegie Mellon University

^dSchool of Information, University of California, Berkeley

Abstract. Motivated attackers cannot always be blocked or deterred. In the physical-world security context, examples include suicide bombers and sexual predators. In computer networks, zero-day exploits unpredictably threaten the information economy and end users. In this paper, we study the conflicting incentives of individuals to act in the light of such threats.

More specifically, in the weakest target game an attacker will *always* be able to compromise the agent (or agents) with the lowest protection level, but will leave all others unscathed. We find the game to exhibit a number of complex phenomena. It does not admit pure Nash equilibria, and when players are heterogeneous in some cases the game does not even admit mixed-strategy equilibria.

Most outcomes from the weakest-target game are far from ideal. In fact, payoffs for most players in any Nash equilibrium are far worse than in the game's social optimum. However, under the rule of a social planner, average security investments are extremely low. The game thus leads to a conflict between pure economic interests, and common social norms that imply that higher levels of security are always desirable.

Key words: Security, Economics, Game Theory, Heterogeneity

1 Introduction

Motivated by observations about widespread and frequent security failures, Hal Varian started a conversation on the role of public goods dilemmas in the reliability and security context [19]. We continued this investigation by analyzing three canonical interdependency scenarios (i.e., weakest-link, average effort and best-shot) in the presence of two investment strategies [7]. Under the assumption of these particularly strong interdependencies, a failure to achieve a common protection goal leads to a compromise of the entire network of agents. For ex-

^{*} This research was partially supported by CyLab at Carnegie Mellon under grant DAAD19-02-1-0389 from the Army Research Office, and by the National Science Foundation under ITR award CCF-0424422 (TRUST).

ample, in the weakest-link game the lack of protection effort by a single agent will immediately be exploited by an attacker to harm all agents.²

However, such strong interdependency effects are not always present in practice or attackers will not be able to exploit them efficiently under all circumstances. Similarly, even with significant protection investments a highly motivated attacker can rarely be fully blocked and deterred.³ Rather, many situations result in asymmetric security consequences, i.e., some agents have their security violated while others remain unharmed.⁴ We considered this case by proposing a novel strategic security interaction: the weakest-target game [7].⁵ Here, an attacker will *always* be able to compromise the agent (or agents) with the lowest protection level, but will leave all others unscathed. Many financially motivated attacks can be explained by considering the weakest target game. For example, botnet herders need to compromise a large number of computing resources at low cost to implement several of their small margin business concepts (e.g., spam distribution [9]). For such purposes, miscreants frequently utilize relatively complex malware that supports a number of attack strategies [18]. A comprehensive defense against such malware becomes increasingly cumbersome because of the inclusion of zero-day exploits, the delivery via different channels (automated scans, email, peer-to-peer networks, ...) and the targeting of different operating systems.⁶

In this paper, we conduct an in-depth investigation of the weakest-target game considering homogeneous and heterogeneous agents. We add to our previous work [7, 8] by deriving more general results and related proofs about properties of the game. In particular, we provide a non-existence proof for pure Nash equilibria, and exact conditions for mixed Nash equilibria for 2-player and N -player games under different parameter conditions. We also discuss important effects that result from the group dynamics inherent in the game.

The weakest target game is not anticipated to be a “bearer of good news” to the security community because it further exacerbates the conflicting incentives of defensive actors. On the one hand, agents, in their avoidance efforts to become the weakest target, may want to engage in significant security investments. On the other hand, the availability of a cheap mitigation alternative weakens the

² Variations can be considered with less strict contribution and associated failure conditions. See, for example, the literature on better-shot and weaker-link games [4].

³ See, for example, the recent successful attacks against Google, Visa/Mastercard, and the US government.

⁴ In complementary work, we also discussed this observation and distinguished more generally between tightly and loosely coupled networks [6].

⁵ We discussed a simplified version of this game initially in the context of network economics [3].

⁶ Consider, for example, the Stuxnet worm that carried four zero-day exploits. Further, it initially infected targets via USB drives, while newer information suggested that Stuxnet also replicated via computer networks. See: <http://www.eweek.com/c/a/Security/Sophisticated-Stuxnet-Worm-Uses-4-Microsoft-Zeroday-Bugs-629672/>.

incentives of all agents to invest in prevention. A similar effect occurs when the population of agents includes at least one subject that has little of value to lose. The result is a game with a particularly perverse set of incentives yielding very ineffective defensive strategies.

In the remainder of this paper, we first conduct a brief discussion on relevant concepts in classical game-theory and security economics. We then present the mathematical model and analysis before offering concluding remarks.

2 Related Work

The strategic aspects apparent in the weakest-target game are complementary to incentive structures analyzed in diverse games in the area of conflict studies. For example, in the Game of Chicken, two agents are driving a car towards a cliff. Both agents can agree to partake in a low-payoff safe choice by stopping early. Alternatively, one of them might dare to deviate by jumping from the car late causing an increase in her own reputation and a reduction of the other’s social status. But if both players opt for the daring choice, the result is of disastrous proportions: aiming for an edge, both agents fall off the cliff [16]. In this prominent example for an anti-coordination game, there are two pure Nash equilibria with asymmetric appeal to the players. However, drivers can agree to disagree by selecting a mixed Nash to moderate between the agents’ desires (e.g., [5]).

In the Prisoner’s Dilemma, rational agents fail to cooperate when facing the enticement of individually beneficial defection from the socially desirable outcome [15]. Following these incentives, the players have to settle for a low-payoff Nash equilibrium. The Prisoner’s Dilemma has been motivation for recent research work on interdependent security in which an agent can invest in protection effort against attacks directly targeting her, but is helpless if an attack is unknowingly spread by her peers [11]. Interdependent security games are helpful to better understand large group effects in the presence of misaligned or conflicting incentives. For example, these games exhibit strong tipping effects that can shift the economy of agents from full prevention efforts to passivity, and vice versa [10]. (See also [12] for an analysis with a similar focus.)

Recent work on linear influence networks also introduce complexity in the decision-making process. Linear influence networks allow for the fine-grained modeling of asset and risk interdependencies (see [13] and [14]). These studies offer an alternative approach to capturing diversity and heterogeneity of incentives for security decision-making. For further research considering the importance of network interdependencies, we also refer to a number of recent review efforts in the area of security economics [1, 2, 17].

3 Model

Each of $N \in \mathbb{N}$ players is responsible for choosing security investments for a single computer that is connected to other computers through a network. The

network is subject to the risk of an external breach that occurs with exogenous probability $p \in (0, 1]$. If the network is breached, the attacker finds the player (or players) with the least amount of protection investment and obliterates her (their) computer(s).⁷ All may not be lost though, as players may choose among two types of security investments to mitigate against damages of a successful breach. They may choose a protection investment, which benefits the public network, and is exemplified by investments such as installing antivirus software or firewalls; or they can choose a *self-insurance* investment, which benefits only the contributing user and is exemplified by an investment such as maintaining extensive private data backups [7]. A full protection investment costs $b_i \in \mathbb{R}^+$ to player i , and a full self-insurance investment costs $c_i \in \mathbb{R}^+$. Players may also choose a partial investment. The choice variables are thus a protection investment level $e_i \in [0, 1]$ and a self-insurance investment level $s_i \in [0, 1]$. Player i begins the game with an initial endowment $M_i \in \mathbb{R}^+$, and suffers a maximum loss of $L_i \in \mathbb{R}^+$ if a security breach occurs. The utility for player i as a result of the investment choice (e_i, s_i) is given by

$$U_i(e_i, s_i) = M_i - pL_i \cdot 1_{e_i \leq \min_{j \neq i} e_j} \cdot (1 - s_i) - be_i - cs_i, \quad (1)$$

where $1_{e_i \leq \min_{j \neq i} e_j} = 1$ if $e_i \leq \min_{j \neq i} e_j$ and 0 otherwise.

4 Analysis

We begin by proving that the game does not admit a pure strategy Nash equilibrium. The result holds for any number of players, and assumes only that the cost and risk parameters (pL_i, c_i, b_i) are all positive.

We next provide a complete characterization of mixed-strategy Nash equilibria in the weakest target game with two players. In brevity, we show that, when either the parameters are symmetric, or when the maximum payoff of players is determined by protection costs rather than self-insurance costs, then exists a well-defined mixed strategy equilibrium. If parameters are not symmetric and self-insurance costs for one player are low, we show that a mixed strategy equilibria does not exist.

Finally, we address the case of N players. We begin by exhibiting a mixed strategy equilibrium in the homogeneous version. Then we derive a framework for addressing the full heterogeneous version.

4.1 No Pure Strategies

Theorem 1. *The weakest target game does not admit a pure strategy Nash equilibrium.*

⁷ In previous work, we referred to this scenario as the weakest target game without mitigation. A slightly different version (with mitigation) allows agents to invest in full protection with the benefit of immunity from attacks [7].

Proof. We will divide into three parts the set of configurations in which each player plays a pure strategy. For each part we will then show that a strategy configuration of the prescribed type fails to be a Nash equilibrium.

First suppose that the strategy configuration has $e_i = 0$ for every i . In this case, the utility of each player i is $M_i - \min\{c_i, pL_i\}$. In such a configuration, if any one player i were to play $(e_i, s_i) = (\epsilon, 0)$ with $\epsilon < \min\{c_i, pL_i\}$, then she would no longer be a weakest target, and her payoff would improve to $M_i - \epsilon > M_i - \min\{c_i, pL_i\}$. Thus, a strategy configuration of this type is not a Nash equilibrium.

Next suppose that the strategy configuration has $e_i = x$ for every i , with $x > 0$. In this case, every player is a weakest target, the utility of player i is $M_i - b_i x - pL_i$. Player i could now improve her utility by playing $(e_i, s_i) = (0, 0)$, and reaping $M_i - pL_i > M_i - b_i x - pL_i$. Thus, such a strategy configuration cannot be a Nash equilibrium.

Finally, suppose the strategy configuration is such that some two players i and j have different protection investment levels (say $e_i < e_j$). In such a configuration, player j is not the weakest target, and thus has a utility of $M - b_j e_j$. If player j were to select a slightly lower investment level, say $x = \frac{e_j + e_i}{2}$, then the corresponding utility would become $M_j - b_j x > M_j - b_j e_j$. Thus, a configuration of this type cannot be a Nash equilibrium.

This exhausts all cases of pure strategies. We have shown that none of the cases is a Nash equilibrium strategy. Thus, no pure strategy Nash equilibrium can exist.

4.2 Mixed Strategies

Mixed strategy descriptions Best response pure strategies for agent i always have one of the two forms: $(e_i, s_i) = (0, 1)$ or $(e_i, s_i) = (x, 0)$ for some $x \in [0, 1]$. So to describe a mixed strategy over this set of pure strategies, it suffices to specify the probability of playing $s_i = 1$ given $e_i = 0$, and a cumulative distribution function $F_i : \mathbb{R} \rightarrow [0, 1]$ defined such that $F_i[x]$ is the probability that $e_i < x$.⁸

We will dispense with giving the first part of the equilibrium conditions (involving self-insurance versus passivity in the case of no protection investment) because that part of the strategy is trivial to determine⁹ and it does not affect whether the strategy is part of an equilibrium.¹⁰ So, to describe a mixed strategy for player i , it suffices to define a non-decreasing left-continuous function $F_i : \mathbb{R} \rightarrow [0, 1]$ satisfying $F_i(x) = 0$ for $x \leq 0$ and $F_i(x) = 1$ for $x > 1$.

⁸ The use of $<$ in our definition of F_i differs from the standard treatment of cumulative distributions, which uses \leq instead. We adopt the former convention so that the probabilities we care about are easy to describe in terms of F_i .

⁹ The structure of the game dictates that player i would self-insure in this instance if and only if $c_i \leq pL_i$

¹⁰ The choice between self-insuring and remaining passive poses no externalities.

Mixed strategy equilibria for 2 players

Theorem 2. *In a two-player weakest target game with parameter conditions satisfying $\frac{\min\{b_1, c_1, pL_1\}}{b_1} = \frac{\min\{b_2, c_2, pL_2\}}{b_2}$, the following mixed strategy is a Nash equilibrium.*

$$F_1(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_2 x}{pL_2} + 1 - \frac{\min\{c_2, pL_2\}}{pL_2} & \text{for } x \in \left(0, \frac{\min\{b_2, c_2, pL_2\}}{b_2}\right] \\ 1 & \text{for } x > \frac{\min\{b_2, c_2, pL_2\}}{b_2} \end{cases}$$

$$F_2(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_1 x}{pL_1} + 1 - \frac{\min\{c_1, pL_1\}}{pL_1} & \text{for } x \in \left(0, \frac{\min\{b_1, c_1, pL_1\}}{b_1}\right] \\ 1 & \text{for } x > \frac{\min\{b_1, c_1, pL_1\}}{b_1} \end{cases}$$

Proof. See Appendix 6.1

Theorem 3. *In a two-player weakest target game with parameter conditions satisfying $\frac{\min\{b_1, c_1, pL_1\}}{b_1} < \frac{\min\{b_2, c_2, pL_2\}}{b_2}$ and $pL_1 \leq c_1$, the following mixed strategy is a Nash equilibrium.*

$$F_1(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_2 x}{pL_2} + 1 - \frac{b_2 L_1}{b_1 L_2} & \text{for } x \in \left(0, \frac{pL_1}{b_1}\right] \\ 1 & \text{for } x > \frac{pL_1}{b_1} \end{cases}$$

$$F_2(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_1 x}{pL_1} & \text{for } x \in \left(0, \frac{pL_1}{b_1}\right] \\ 1 & \text{for } x > \frac{pL_1}{b_1} \end{cases}$$

Proof. See Appendix 6.2

Theorem 4. *In the remaining parameter condition, $\frac{\min\{b_1, c_1, pL_1\}}{b_1} < \frac{\min\{b_2, c_2, pL_2\}}{b_2}$ and $c_1 < pL_1$, there does not exist a mixed-strategy Nash equilibrium.*

Proof. See Appendix 6.3

Mixed strategy equilibria for N players First we consider what happens in a game where the parameters are homogeneous – that is, $c_i = c$, $b_i = b$, and $L_i = L$. This scenario was addressed for a limited parameter range in [7].

Theorem 5. *In the case of homogeneous parameters, a Nash equilibrium occurs when each player plays the following mixed strategy.*

$$F_i(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ 1 - \sqrt[N-1]{\frac{\min\{c, pL\} - bx}{pL}} & \text{for } x \in \left(0, \frac{\min\{b, c, pL\}}{b}\right] \\ 1 & \text{for } x > \frac{\min\{b, c, pL\}}{b} \end{cases}$$

Proof. See appendix 6.4

With fully heterogeneous parameters, the analysis becomes significantly more complex. We derive a basic framework for a mixed equilibrium strategy involving N heterogeneous agents, but omit the exact specification of the necessary conditions for space reasons and will include them in the extended version.

Theorem 6. *The following mixed strategy describes an equilibrium in which each player j receives an expected utility of $M_j - \min\{c_j, pL_j\}$ for a continuous range of plays. (Assume WLOG that $\frac{\min\{b_1, c_1, pL_1\}}{b_1} \leq \frac{\min\{b_j, c_j, pL_j\}}{b_j}$ for $j > 1$).*

$$F_i(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ 1 - \frac{N-1 \sqrt{\prod_{j=1}^N \frac{\min\{c_j, pL_j\} - b_j x}{pL_j}}}{\frac{\min\{c_i, pL_i\} - b_i x}{pL_i}} & \text{for } x \in \left(0, \frac{\min\{b_1, c_1, pL_1\}}{b_1}\right] \\ 1 & \text{for } x > \frac{\min\{b_1, c_1, pL_1\}}{b_1} \end{cases}$$

Proof. See Appendix 6.5

4.3 Social Optimum: The sacrificial lamb

In the weakest target game, the least protected players always bear the brunt of the attack. The socially desirable outcome then ensures a minimization of the aggregate loss for all players. Typically, a planner will elect a specific agent, i.e., the sacrificial lamb, to invest in the smallest possible security effort to attract the attacker. If self-insurance is less costly than the loss from a security compromise, then the planner will invest in mitigation for the designated agent. Under heterogeneity assumptions, the social planner must undertake this comparison across all agents, i.e., she typically needs to identify the two agents who have the least to lose, $c_{min} = \min\{c_i\}$, and the lowest cost self-insurance option, $pL_{min} = \min\{pL_j\}$, respectively.¹¹ She will then select the agent with the lower of the two values. All other agents will merely invest in a token security effort, ϵ , to escape the attack.

With this strategy configuration, the total cost for all players is approximately $\min\{c_{min}, pL_{min}\} + \epsilon$. This is the best strategy possible up to a factor of epsilon.

5 Discussion

In the absence of the mediating presence of a social planner, agents are struggling to find a cheap way to ameliorate the threat from the attacker and the behavior of the agents can become quite complex. In the following, we illustrate two main behavioral trends related to the mixed strategy play identified in the analysis section.

First, when for all agents the cost of protection is lower than the cost of self-insurance, $b_i < c_i$, and preventive efforts are considered worthwhile, $b_i < pL_i$,

¹¹ It is, of course, possible that a single agent acts in both roles, $i = j$.

then agents' behavior can be described as a "race against the wall," i.e, they try to avoid a security compromise by selecting a very high protection effort. But according to the rules of the game an escape from the attacker's wrath is not possible and agents understand that collusion on the highest protection level would only benefit the attacker and would be wasteful from a cost perspective. Therefore, all agents probabilistically lower their security efforts to a certain degree with smaller probabilities assigned to lower protection levels. That is, agents engage in an implicit and tacit process of risk sharing. Theorem 2, when $b_i < \min\{c_i, pL_i\}$, is an example for this type of behavior.

Second, a more nuanced behavior can be observed when the population includes at least one player with a low self-insurance cost, $c_i < b_i$ and $c_i < pL_i$, or a low potential loss, $pL_i < b_i$. We term this agent an implicit leader in the game. It is easy to see that she would not invest more than a certain threshold amount, $e_{max} = \frac{\min\{c_i, pL_i\}}{b_i}$, in protection efforts. And other players can infer that they will lose at most $b_j \cdot \left(\frac{\min\{c_i, pL_i\}}{b_i} + \epsilon\right)$ by investing a small amount more than e_{max} in preventive efforts. However, to achieve an equilibrium outcome close to this scenario the other players need to motivate the leader to invest in e_{max} , otherwise, the game play would unravel. That is, they need to ensure that for a whole range of parameters, the leader has an incentive to invest in protection efforts. In order to achieve this goal they need to "support the leader and share the burden" by engaging in probabilistic protection efforts below e_{max} . An example for this scenario is put forward in Theorem 3.

Adding more players and more heterogeneous preferences to the game play increases the complexity of the outcome. For example, if an agent is endowed with a extremely low cost of self-insurance it is very difficult to persuade her to act as a leader. In future work we will more thoroughly cover the nuances in the N -player version of the game, and illustrate the important behaviors with graphical representations. In the future, we also plan to engage in experimental validations of the predicted outcomes.

6 Conclusion

The weakest target game is interesting and well-motivated, but difficult to analyze. It does not admit pure Nash equilibria, and when players are heterogeneous in some cases the game does not even admit mixed-strategy equilibria. When mixed strategy equilibria do occur they are dominated by phenomena which we identified in the discussion section.

Most outcomes are far from ideal. In fact, *payoffs* for most players in any Nash equilibrium are far worse than in the game's social optimum. However, under the social planner rule *average security investments* are extremely low. This leads to a conflict between budgetary interests and a desire for increased security readiness.

References

1. R. Anderson and T. Moore. The economics of information security. *Science*, 314(5799):610–613, October 2006.
2. R. Böhme and G. Schwartz. Modeling cyber-insurance: Towards a unifying framework. In *Proceedings of the Ninth Workshop on the Economics of Information Security (WEIS'10)*, Cambridge, MA, June 2010.
3. N. Christin, J. Grossklags, and J. Chuang. Near rationality and competitive equilibria in networked systems. In *Proceedings of ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, pages 213–219, Portland, OR, August 2004.
4. R. Cornes and T. Sandler. *The theory of externalities, public goods, and club goods*. Cambridge University Press, Cambridge, UK, 1986.
5. A. Dixit and S. Skeath. *Games of Strategy*. Norton & Company, New York, NY, 1999.
6. N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In *Proceedings of the 13th International Conference Financial Cryptography and Data Security (FC'09), Lecture Notes in Computer Science (LNCS), No. 5628*, pages 167–183, Christ Church, Barbados, February 2009.
7. J. Grossklags, N. Christin, and J. Chuang. Secure or insure? A game-theoretic analysis of information security games. In *Proceedings of the 2008 World Wide Web Conference (WWW'08)*, pages 209–218, Beijing, China, April 2008.
8. J. Grossklags, N. Christin, and J. Chuang. Security and insurance management in networks with heterogeneous agents. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*, pages 160–169, Chicago, IL, July 2008.
9. C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the Conference on Computer and Communications Security (CCS)*, pages 3–14, Alexandria, VA, October 2008.
10. M. Kearns and L. Ortiz. Algorithms for interdependent security games. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 561–568. MIT Press, Cambridge, MA, 2004.
11. H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2–3):231–249, March 2003.
12. M. Lelarge and J. Bolot. Network externalities and the deployment of security features and protocols in the Internet. In *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS'08)*, pages 37–48, Annapolis, MA, June 2008.
13. A. Miura-Ko, B. Yolken, J. Mitchell, and N. Bambos. Security decision-making among interdependent organizations. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 66–80, Pittsburgh, PA, June 2008.
14. K. Nguyen, T. Alpcan, and T. Basar. Stochastic games for security in networks with interdependent nodes. In *Proceedings of the International Conference on Game Theory for Networks (GameNets 2009)*, pages 697–703, Istanbul, Turkey, May 2009.
15. A. Rapoport and A. Chammah. *Prisoner's Dilemma: A Study in Conflict and Cooperation*. Ann Arbor Paperbacks, University of Michigan Press, 1965.
16. A. Rapoport and A. Chammah. The game of chicken. *American Behavioral Scientist*, 10(3):10–28, November 1966.

17. S. Roy, C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya, and Q. Wu. A survey of game theory as applied to network security. In *Proceedings of the 43rd Hawaii International Conference on System Sciences (HICSS'10)*, pages 1–10, Koloa, HI, January 2010.
18. E. Skoudis. *Malware: Fighting malicious code*. Prentice Hall, Upper Saddle River, NJ, 2004.
19. H. Varian. System reliability and free riding. In J. Camp and S. Lewis, editors, *Economics of Information Security (Advances in Information Security, Volume 12)*, pages 1–15. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

Appendix

6.1 Proof of Theorem 2

Proof. First note that under the specified parameter conditions, each F_i is a left-continuous increasing function on \mathbb{R} , that $F_i(x) = 0$ for $x \leq 0$ and that $F_i(x) = 1$ for $x \geq 1$. Thus F_i describes a valid mixed strategy for player i . To show that the mixed-strategy configuration is a Nash equilibrium, we will consider two parameter sub-cases separately.

Case 1: $b_i \leq \min\{c_i, pL_i\}$ for $i = 1, 2$

In this case, we have $\frac{\min\{b_1, c_1, pL_1\}}{b_1} = 1$. So the mixed strategies simplify to:

$$F_1(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_2x}{pL_2} + 1 - \frac{\min\{c_2, pL_2\}}{pL_2} & \text{for } x \in (0, 1] \\ 1 & \text{for } x > 1 \end{cases}$$

$$F_2(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_1x}{pL_1} + 1 - \frac{\min\{c_1, pL_1\}}{pL_1} & \text{for } x \in (0, 1] \\ 1 & \text{for } x > 1 \end{cases}$$

As the strategies are now symmetric it suffices to assume that player 1 is playing F_1 and show that player 2's response strategy is optimal. So assume player 1 is playing F_1 .

- If player 2 plays $e_2 = 0$, her resulting payoff will be $M_2 - \min\{c_2, pL_2\}$.
- If player 2 plays $e_2 = x \in (0, 1]$, her resulting payoff will be $M_2 - b_2x - pL_2(1 - F_1(x)) = M_2 - b_2x - pL_2 \left(\frac{\min\{c_2, pL_2\}}{pL_2} - \frac{b_2x}{pL_2} \right) = M_2 - \min\{c_2, pL_2\}$

We see that player 2 receives the same payoff regardless of her choice of strategies. Thus, playing a mixed strategy distribution over all possible strategies is an optimal response strategy, and hence the strategy configuration is a mixed-strategy equilibrium.

Case 2: $\min\{c_1, pL_1\} < b_1$ and $\frac{\min\{c_1, pL_1\}}{b_1} = \frac{\min\{c_2, pL_2\}}{b_2}$

In this case the mixed strategies simplify to:

$$F_1(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_2x}{pL_2} + 1 - \frac{\min\{c_2, pL_2\}}{pL_2} & \text{for } x \in \left(0, \frac{\min\{c_2, pL_2\}}{b_2}\right] \\ 1 & \text{for } x > \frac{\min\{c_2, pL_2\}}{b_2} \end{cases}$$

$$F_2(x) = \begin{cases} 0 & \text{for } x \leq 0 \\ \frac{b_1 x}{pL_1} + 1 - \frac{\min\{c_1, pL_1\}}{pL_1} & \text{for } x \in \left(0, \frac{\min\{c_1, pL_1\}}{b_1}\right] \\ 1 & \text{for } x > \frac{\min\{c_1, pL_1\}}{b_1} \end{cases}$$

Again the strategies are symmetric so it suffices to assume player 1 is playing F_1 and consider the best response of player 2.

- Again, if player 2 plays $e_2 = 0$, her resulting payoff will be $M_2 - \min\{c_2, pL_2\}$.
- If player 2 plays $e_2 = x \in \left(0, \frac{\min\{c_2, pL_2\}}{b_2}\right]$, her resulting payoff will be $M_2 - b_2 x - pL_2(1 - F_1(x)) = M_2 - b_2 x - pL_2 \left(\frac{\min\{c_2, pL_2\}}{pL_2} - \frac{b_2 x}{pL_2}\right) = M_2 - \min\{c_2, pL_2\}$.
- If player 2 were to play $e_2 = x > \frac{\min\{c_2, pL_2\}}{b_2}$, her resulting payoff would be $M_2 - b_2 x - pL_2(1 - F_1(x)) = M_2 - b_2 x < M_2 - b_2 \cdot \frac{\min\{c_2, pL_2\}}{b_2} = M - \min\{c_2, pL_2\}$.

We see that the first two options yield the same payoff, and the third option yields a suboptimal payoff. Since response strategy for player 2 described by F_2 is a mixed strategy over pure strategies of only the first two forms, it is a best response strategy. We see again in this case that the mixed strategy configuration is a Nash equilibrium.

6.2 Proof of Theorem 3

Proof. Again F_i describes a valid mixed strategy for player i . Assume that player 1 is playing F_1 and consider the utility of player 2's response strategy.

- If player 2 were to play $e_2 = 0$, she would reap $M_2 - \min\{c_2, pL_2\}$. (Note that according to F_2 she plays this strategy with probability zero.)
- If player 2 plays $e_2 = x \in \left(0, \frac{pL_1}{b_1}\right]$, her payoff is $M_2 - b_2 x - pL_2(1 - F_1(x)) = M_2 - b_2 x - pL_2 \left(\frac{b_2 L_1}{b_1 L_2} - \frac{b_2 x}{pL_2}\right) = M_2 - \frac{b_2 L_1}{b_1} > M_2 - b_2 \cdot \frac{\min\{c_2, pL_2\}}{b_2} = M - \min\{c_2, pL_2\}$.
- Finally, if player 2 were to play $e_2 = x > \frac{pL_1}{b_1}$, her payoff would be $M_2 - b_2 x < M_2 - b_2 \cdot \frac{pL_1}{b_1}$.

We see that the optimal payoff player 2 can achieve is $M_2 - \frac{b_2 pL_1}{b_1}$. She achieves this utility exactly when she plays $x \in \left(0, \frac{pL_1}{b_1}\right]$; and this is exactly the set of strategies that she plays according to her mixed-strategy specification F_2 .

Next assume that player 2 is playing F_2 and consider the utility of player 1's response strategy.

- If player 1 plays $e_1 = 0$, she reaps $M - pL_1$.
- If player 1 plays $e_1 = x \in \left(0, \frac{pL_1}{b_1}\right]$, she reaps $M_1 - b_1 x - pL_1(1 - F_2(x)) = M_1 - b_1 x - pL_1 \left(1 - \frac{b_1 x}{pL_1}\right) = M - pL_1$.

– Lastly, if player 1 plays $e_1 = x > \frac{pL_1}{b_1}$, she reaps $M_1 - b_1x < M_1 - b_1 \cdot \frac{pL_1}{b_1} = M_1 - pL_1$.

We see that player 1 maximizes her utility by playing any of the first two strategy conditions, and this conforms to the prescription of F_1 . So player 1 is playing an optimal response strategy.

This completes the proof that this strategy configuration is a Nash equilibrium.

6.3 Proof of Theorem 4

Proof. First note that the parameter conditions imply that $\frac{\min\{b_1, c_1, pL_1\}}{b_1} \neq 1$ and hence $c_1 < \min\{pL_1, b_1\}$.

Suppose that there does exist a mixed strategy Nash equilibrium under these parameter conditions. For $i = 1, 2$, let X_i be the set of pure strategies that occur in player i 's mixed strategy; and let F_i be the cumulative distribution function for e_i in player i 's mixed strategy, defined so that $F_i(x) = Pr[e_i < x]$. Note that each F_j is monotone non-decreasing and left-continuous.¹²

We next prove a sequence of lemmas that give more structure to the functions F_j . Ultimately, these lemmas will result in a contradiction involving the behavior of F_2 near the point $x = 0$, demonstrating that functions satisfying the prescribed properties cannot exist.

Lemma 1. *There exists a real number β with $0 < \beta \leq \frac{c_1}{b_1}$ such that for each j , $F_j(\beta) = 1$, but for every real number $\alpha < \beta$, $F_j(\alpha) < 1$.*

Proof. The utility of player 1's mixed strategy is at least $M_1 - c_1$. Hence any choice of $e_1 > \frac{c_1}{b_1}$ is deterministically suboptimal. Let β be the least upper bound on e_1 in player 1's mixed strategy. (More formally, we could define β to be the maximum element in $\overline{X_1}$). Then $\beta \leq \frac{c_1}{b_1}$. Interpreting the definition of least upper bound into the language of F_1 , we also have $F_1(\alpha) < 1$ for every $\alpha < \beta$, and $F_1(\gamma) = 1$ for every $\gamma > \beta$.

Now, player 2 must have elements in her mixed strategy that take e_2 arbitrarily close to β from below. Otherwise, player 1 would have chosen an upper bound lower than β to obtain a better utility. Thus we have $F_2(\alpha) < 1$ for every $\alpha < \beta$. Also player 2 cannot have any part of her mixed strategy include a protection level strictly higher than β , for otherwise she would have preferred to reduce this expenditure by a small amount to be closer to β and obtain an improved utility. Thus $F_2(\gamma) = 1$ for every $\gamma > \beta$.

Next we see that player 1 cannot play the pure strategy $e_1 = \beta$ with positive probability. Otherwise, the discontinuity of F_1 at β would cause player 2 to

¹² Note that our use of $<$ as opposed to \leq differs from the standard treatment of cumulative distribution functions. Our notation yields left-continuous as opposed to right continuous. The reason we use the formulation with $<$ is that we need to know when $e_i \geq x$; this is easily expressed algebraically in terms of the predicate $e_i < x$, but using \leq would make it cumbersome.

receive a strictly higher payoff from playing $\beta + \epsilon$ (for sufficiently small ϵ) than from playing $\beta - \epsilon$. This contradicts the presumed optimality of player 2's mixed strategy, which contains plays greater than $\beta - \epsilon$ for every ϵ but no plays of $\beta + \epsilon$. Similarly $F_2(\beta) = 1$ by a completely analogous argument to the one above.

This completes the proof of the lemma.

Lemma 2. $\lim_{x \rightarrow 0^+} F_2(x) = 0$

Proof. In words, this lemma says that player 2's mixed strategy cannot contain a pure strategy component of the form $e_2 = 0$ with positive probability. To see this, observe that the maximum benefit player 2 can achieve from playing $e_2 = 0$ is $M_2 - \min\{c_2, pL_2\}$. However, using the result from the previous lemma, if player 2 were to play $e_2 = \frac{c_1}{b_1}$, she would fail to be the weakest target with probability 1, and would thus receive a utility of $M - b_2 \cdot \frac{c_1}{b_1}$. Since our parameter conditions imply $\frac{c_1}{b_1} < \frac{\min\{c_2, pL_2\}}{b_2}$, we have $M - b_2 \cdot \frac{c_1}{b_1} > M - b_2 \cdot \frac{\min\{c_2, pL_2\}}{b_2} = M - \min\{c_2, pL_2\}$. Thus playing $e_2 = 0$ is a suboptimal strategy for player 2. I.e., $\lim_{x \rightarrow 0^+} F_2(x) = 0$.

Lemma 3. For $j \neq i$, and for $w, x \in X_j$, we have $F_j(w) = F_j(x) - \frac{b_i}{pL_i}(x - w)$.

Proof. From the weakest target game definition, player j loses pL_j whenever $0 < e_j \leq e_i$, and this happens with probability $1 - F_i(e_j)$. We see that for $x > 0$ and for $j \neq i$ the utility of player j is directly related F_i via

$$U_j(x, 0) = M_j - b_j x - pL_j(1 - F_i(x)).$$

Now in a mixed strategy equilibrium all pure strategy components yield the same utility, hence for each $w, x \in X_i$ we have $U_i(w) = U_i(x)$. By rewriting the expression in terms of F_j , we obtain the result

$$F_j(w) = F_j(x) - \frac{b_i}{pL_i}(x - w).$$

Lemma 4. For $j \neq i$, and for $w \in X_j$, we have $F_j(w) = 1 - \frac{b_i}{pL_i}(\beta - w)$.

Proof. Since each F_j is left continuous, we have $\lim_{\alpha \rightarrow \beta^-} F_j(\alpha) = F_j(\beta) = 1$. Let $w \in X_j$ and let $\langle \alpha_n \rangle_{n \in \mathbb{N}}$ be a sequence from X_j that converges to β . (Such a sequence exists in X_j from the arguments given in Lemma 1.) Then from Lemma 3 we have $F_j(w) = F_j(\alpha_n) - \frac{b_i}{pL_i}(\alpha_n - w)$ for each n . Taking the limit of both sides yields

$$F_j(w) = 1 - \frac{b_i}{pL_i}(\beta - w).$$

Lemma 5. For $j \neq i$, and for $z \in (0, \beta)$, we have $F_j(z) \geq 1 - \frac{b_i}{pL_i}(\beta - z)$.

Proof. If $z \notin X_i$, then we cannot use equality of utilities, but using Lemma 4 and the properties of F_j as a cumulative distribution, we have for $z \in (0, \beta)$:

$$\begin{aligned}
F_j(z) &= \inf\{F_j(w) : w \in X_j \text{ and } w > z\} \\
&= \inf\left\{1 - \frac{b_i}{L_i}(\beta - w) : w \in X \text{ and } w > z\right\} \\
&\geq 1 - \frac{b_i}{pL_i}(\beta - z)
\end{aligned}$$

Finally, for the punchline,

Using Lemma 5, for every $\epsilon \in (0, \beta)$ we have

$$\begin{aligned}
F_2(\epsilon) &\geq 1 - \frac{b_1}{pL_1}(\beta - \epsilon) \\
&= 1 + \frac{b_1}{pL_1}\epsilon - \frac{b_1}{pL_1}\beta \\
&> 1 - \frac{b_1}{pL_1} \cdot \frac{c_1}{b_1} \\
&= 1 - \frac{c_1}{pL_1}
\end{aligned}$$

In particular, $\lim_{\epsilon \rightarrow 0^+} F_2(\epsilon) \geq 1 - \frac{c_1}{pL_1} > 0$. This contradicts the conclusion of Lemma 2.

We conclude that no mixed strategy equilibrium can exist.

6.4 Proof of Theorem 5

Proof. F_i describes a valid mixed strategy for player i . Assume that all players $j \neq i$ are playing F_j and consider the best response of player i .

- If player i plays $e_i = 0$, she reaps $M - \min\{c, pL\}$.
- If player i plays $e_i = x \in \left(0, \frac{\min\{b, c, pL\}}{b}\right]$, her payoff is $M - bx - pL \cdot \prod_{j \neq i} (1 - F_j(x)) = M - bx - pL \cdot \prod_{j \neq i} \left(1 - \sqrt[N-1]{\frac{\min\{c, pL\} - bx}{pL}}\right) = M - bx - pL \cdot \frac{\min\{c, pL\} - bx}{pL} = M - \min\{c, pL\}$.
- Finally, if player i were to play $e_i = x > \frac{\min\{b, c, pL\}}{b}$, then this is possible only if $\min\{b, c, pL\} = \min\{c, pL\}$, and in this case, her payoff would be at most $M - bx < M - b \cdot \frac{\min\{c, pL\}}{b} = M - \min\{c, pL\}$.

We see that the optimal payoff player i can achieve is $M - \min\{c, pL\}$. She achieves this utility exactly when she plays $x \in \left[0, \frac{\min\{b, c, pL\}}{b}\right]$; and this is exactly the set of strategies that she plays according to her mixed-strategy specification F_i .

6.5 Proof of Theorem 6

Proof. The algebraic part of the formula is derived as follows:

$$\begin{aligned}
M_i - \min\{c_i, pL_i\} &= M_i - b_i x - pL_i \prod_{j \neq i} (1 - F_j(x)) \\
\frac{\min\{c_i, pL_i\} - b_i x}{pL_i} &= \prod_{j \neq i} (1 - F_j(x)) \\
\prod_{k=1}^N \frac{\min\{c_k, pL_k\} - b_k x}{pL_k} &= \prod_{k=1}^N (1 - F_k(x))^{N-1} \\
\frac{\prod_{k=1}^N \frac{\min\{c_k, pL_k\} - b_k x}{pL_k}}{\left(\frac{\min\{c_i, pL_i\} - b_i x}{pL_i}\right)^{N-1}} &= \frac{\prod_{k=1}^N (1 - F_k(x))^{N-1}}{\left(\prod_{j \neq i} (1 - F_j(x))\right)^{N-1}} \\
\frac{\prod_{k=1}^N \frac{\min\{c_k, pL_k\} - b_k x}{pL_k}}{\left(\frac{\min\{c_i, pL_i\} - b_i x}{pL_i}\right)^{N-1}} &= (1 - F_i(x))^{N-1} \\
\sqrt[N-1]{\frac{\prod_{k=1}^N \frac{\min\{c_k, pL_k\} - b_k x}{pL_k}}{\left(\frac{\min\{c_i, pL_i\} - b_i x}{pL_i}\right)^{N-1}}} &= (1 - F_i(x)) \\
1 - \sqrt[N-1]{\frac{\prod_{k=1}^N \frac{\min\{c_k, pL_k\} - b_k x}{pL_k}}{\left(\frac{\min\{c_i, pL_i\} - b_i x}{pL_i}\right)^{N-1}}} &= F_i(x)
\end{aligned}$$