

The cost of misinformation in deadly conflicts

Hawk-Dove games and suicidal terrorism

James E. Hanley, Ph.D.

Visiting Scholar, Workshop in Political Theory and Policy Analysis,
Indiana University, Bloomington, IN 47408
james@jhanley.net

John Orbell, Ph.D.

Political Science Department & Institute of Cognitive and Decision Sciences,
University of Oregon, Eugene, OR 97403
jorbell@oregon.uoregon.edu

Tomonori Morikawa, Ph.D.

Center for International Education, Waseda University,
1-7-14-404, Nishi-Waseda, Shinjuku-ku, Tokyo, 169-0051, Japan
goducks1221@hotmail.com

ABSTRACT. The hijacking and purposeful crashing of airplanes into the World Trade Center and the Pentagon on September 11, 2001, prompts questions about why the passengers and crew of those airplanes did not act to prevent these attacks, as did at least some passengers on a hijacked flight that crashed in Pennsylvania. We argue, first, that humans have an evolved cognitive bias that leads to the expectation that antagonists hope to survive conflict and, second, that highly credible information to the contrary is needed to overcome this bias. Absent such information, the passengers on at least two airplanes incorrectly interpreted the game being played as a hawk-dove version of a conflict-of-interest game, when it was actually a “suicidal terrorism” variant of that family. Given that other terrorists may have been in the air and ready to act, the airlines’ policy of not informing passengers about such events could have risked disabling them from reacting forcefully when force alone was advisable.

Humans live in an ecology of games,¹ a range of potential interactions with other persons whose interests are seldom completely in accord with our own, and whose actions and intentions are often difficult to predict. Consequently, evolution has equipped us with “social” or “political” intelligence, which enables us successfully to navigate the complex space of human interactions.^{2, 3, 4} Among the difficulties we face is simply recognizing what game another person is offering us. This is a crucial task, because we choose our strategies based on the game we think we are playing, and if we are actually caught up in a different game our strategic choices are unlikely to be successful. However, little is yet known

about how we accomplish this initial task, game recognition.⁵

The events of September 11, 2001, show just how costly such recognition errors can be. On that day, terrorists hijacked four American airplanes — and may have intended to hijack more — and deliberately crashed them into buildings filled with people. Two airplanes were flown into the twin towers of the World Trade Center in New York City, causing their collapse and killing over three thousand people, while another airplane hit the Pentagon, killing about two hundred more. The fourth airplane came down in a field near Pittsburgh, Pennsylvania, after passengers became aware of the hijackers’ intentions and fought back.

The cost of misinformation in deadly conflicts

While unsuccessful in winning control, the passengers did still thwart the hijackers' effort to use the airplane itself as a weapon of terror.

Given that humans are equipped with evolved social intelligence, with which they interpret each other's goals and intentions, why did the passengers of two and possibly three of these airplanes apparently not recognize what the hijackers were up to, at least not soon enough and generally enough to act effectively? On one level, of course, the passengers were simply following the standard policy of the airlines, which was to refrain from any actions that might antagonize the hijackers. Closely related to this was a policy of keeping passengers on flights under-way ignorant of such events. However, these policies have been based on a type of hijacking game that is different from the one the September 11 terrorists were playing.

We suspect that our evolved social intelligence contains a bias against recognizing certain types of games, which bias could be righted only by the application of relevant information. The passengers of the airplanes crashed in New York City did not have such information, and the policy that denied it to them was most likely shaped innocently by that same cognitive bias. The passengers made rational choices based on the type of game they reasonably believed the terrorists had forced them into.

Elsewhere in this issue of *PLS*, we develop a theoretical argument about the evolution of human cognition in the context of hawk-dove conflict-of-interest games. Such games are characterized by a conflict of interest over some resource, whose value is designated V . The individuals can either challenge for control of the resource by offering "hawk," signaling a willingness to fight for it, or they can play "dove," signaling their unwillingness to fight. If both play hawk, the winner of the conflict captures V in its entirety while the loser suffers some cost, C . If one plays hawk while the other plays dove, the one playing hawk captures V unopposed and avoids the risk of incurring C , while the one who plays dove neither gains nor benefits. Because the hawk-dove game has no dominating strategy — a strategy that is best regardless of the other person's choice — a rational player's strategic choice is determined by the game's payoffs, V and C , and by the probability of winning a hawk-hawk fight, $p(win)$. The expected value (EV) of playing hawk is therefore:

$$p(win)(V) - [1-p(win)](C).$$

If the result is positive, then the individual should play hawk, but if it is negative, he should play dove.

That analysis made the simplifying assumptions that V and C are known by both players and that they are the same for both players. These assumptions were justified by, respectively, an interest in the evolution of *social* intelligence rather than technical intelligence and the long-term character of evolution, in which payoffs to players should average out across innumerable interactions. Nonetheless, it is particular real-world interactions which are the specifics from which, in aggregate, natural selection works, and in any particular confrontation V and C may *not* be the same for each player. Social intelligence is relevant in such cases, as it helps us perceive in the immediate psychological or physiological state of another player subjective factors that may be crucial in his or her definition of V — more crucial perhaps than some, any, or all objective factors.

Focusing on game recognition on September 11, 2001, we ask three questions:

1. Were passengers able to recognize the game and its parameters, and if not why not?
2. What were the consequences of success or failure in recognition?
3. What can we learn from the outcomes of these desperately serious "games"?

The available evidence suggests that the passengers on the flight that crashed in Pennsylvania responded differently than the passengers on the two flights that crashed in New York. We also know that the passengers on the Pennsylvania plane had access to different information than did the passengers on those other two planes. Accordingly, we begin by analyzing and comparing these differing situations.

It appears that the passengers on the airplanes that crashed into the World Trade Center misinterpreted the parameters of, and hence misidentified, the "game" the terrorists had thrust upon them. Passengers most likely saw the encounter as a fairly standard hawk-dove game, whose parameters consisted of a low $p(win)$ if they chose to fight the hijackers, as well as a relatively low V — control of a pilotless airplane — and a very high C , being killed if they fought and lost. The estimated return for playing dove (passive obedi-

ence to the hijackers) differed from the standard formulation in that it imposed a negative return, fear and inconvenience, but this was nearly negligible compared to *C*. Under such parameters the *EV* of responding to the terrorists' initial hawk choice with a hawk choice of their own was negative, and so they chose to play dove. Given the information available to them, this choice was quite rational: they had *no* reason to suppose that playing dove would result in certain death for everyone aboard the airplane — the hijackers as well as themselves — plus many more people on the ground.

However, the hijackers had in fact forced upon them a significantly different game, which somewhat uncomfortably we call the “suicidal terrorist” game. This also was a conflict-of-interest game and offered the same $p(\textit{win})$ as hawk-dove. But suicidal terrorism presented to the passengers a higher *V*, since they might save many fellow citizens and lessen American losses by trying, even unsuccessfully, to save themselves. Suicidal terrorist also presented a much higher *C*: death for themselves plus failure to save anyone or anything else. Strangely by hawk-dove standards, suicidal terrorist projected a return for playing dove that was *identical* to *C*. While dove can be a rational response to hawk in the hawk-dove game, it cannot be rational in the suicidal terrorist game, if recognized. Once the other player has initiated a suicidal terrorist game with an aggressive move, fighting back is the dominating strategy because it is the only choice that offers some probability, however slight, of avoiding *C*.

By retrospectively analyzing their apparent strategic choice, then, it appears that the passengers were *not* able to recognize the game and its parameters. This leads us to the second part of our first question: why did they not?

Note that the passengers did correctly identify the game as a conflict-of-interest game, meaning they did recognize the game's basic structure. What they misinterpreted was the terrorists' subjective *V* and *C*, which were different from what they would have been in a hawk-dove game. Whereas normally a cost of losing, for the terrorists *death was actually part of the value they were seeking* — more specifically, the rewards of paradise that dying in pursuit of their mission would bring them. This eagerness to incur what is normally considered the ultimate cost was a parameter understandably difficult for the passengers to interpret.

First, in interpreting the terrorists' subjective *V* and *C*, the passengers relied on the best available information. Most immediately relevant was their knowledge of what had happened in previous hijackings, drawn from their memory of media reports. Since the great majority of previous hijackings had *not* been committed by people who intended their own deaths, this “best available” information would not have helped the passengers infer that these particular hijackers did, in fact, intend to die. As well, these remembered news accounts would have informed the passengers that most previous victims had survived as a result of *not* attacking their captors, thus supporting their interpretation that the return for playing dove, while negative, was far better than *C*.

In conjunction with this, at least some of the passengers would have known or would have learned from the actions of the flight crew that airline policy for dealing with hijackers was designed to minimize the risk of physical harm to the innocent. This policy obviously assumed that hijackers would hope to survive, and it would have authoritatively reinforced a hawk-dove interpretation of the immediate confrontation.

Within this context, the hijackers in at least one of these two airplanes purposefully gave the passengers signals suggesting that they were engaged in an ordinary hijacking, telling them “just stay quiet, and you'll be okay”; they were going to return to their airport of origin, and “if you make any moves you'll just endanger yourself and the airplane.”⁶ Our social intelligence is designed in part for the task of interpreting such signals, searching for whatever clues may help differentiate accurate information from manipulative messages.⁷ One of those clues is context, and in this case context supported the unfortunately erroneous conclusion that the hijackers' signals were accurate information.

But this does not fully explain why the passengers did not perceive that these *particular* hijackers were different, and that something in their mental state created non-standard game parameters. Nor is it more than a proximate explanation for why the airlines had policies of avoiding conflict and withholding information. More fundamental to the passengers' decision-making that day was the fact that we do not expect people who attack us to put their own death in the value column. The likely reason for this is an evolved cognitive bias that leads us to predict that our antagonist at least *hopes* to survive.

The cost of misinformation in deadly conflicts

As we argued in our other article, conflict-of-interest games must have been and must still be an important part of our adaptive environment. Our particular reference, however, was to hawk-dove games, which are common in animal encounters, from the mating contests that inspired the initial formalization of the game⁸ all the way up to war between nations. Not only humans but most of their ancestors would have engaged in hawk-dove games on a regular basis. Thus we predicted that humans have multiple levels of evolved cognitive capacities for playing them.⁹

Granted, suicidal terrorism games are also not unknown. The medieval era had the Islamic Cult of the Assassins and World War II the *kamikaze*; Palestinian insurgencies in Israel recruit suicide bombers and the Tamil insurgency in Sri Lanka the Black Tigers. Yet these instances stand out because they are unusual. Among conflict-of-interest games, hawk-dove is commonplace, but suicidal terrorism is extremely rare.

Suicidal terrorism is rare because the successful suicidal terrorist is limited in his ability to translate his success into reproductive fitness. Whereas the winner of a hawk-dove game has captured resources that may enhance attractiveness for mating or may improve provision for offspring, the suicidal terrorist can no longer mate or provide. His only path to reproductive success is through kin selection, enhancing his family's reproductive success through his own sacrifice.¹⁰ Appropriate incentives can, of course, be employed consciously. The Palestinian organization, Hamas, for example, promises rich rewards for the families of suicide bombers and has created a recruitment video showing one martyr's parents enjoying their new wealth.¹¹ But the necessity of such consciously employed incentives, as well as the conscious use of eternal paradise as an incentive, underscores the fact that suicidal terrorism is not a game we play casually.

Given the scarcity of suicidal terrorism games and the ubiquity of hawk-dove games, a standing assumption that an opponent wants to survive a conflict-of-interest game must have been adaptive. Thus, rather than having evolved cognitive capacities for recognizing and playing suicidal terrorism games, we are likely to have a cognitive bias that leads us to interpret an identified conflict-of-interest game as being of the standard hawk-dove form. Highly credible information might be required to overcome this bias. The passengers on these

two airplanes did not have such information available, nor was there any way for them to gain it.

In contrast, the passengers of the airplane that crashed in Pennsylvania appear to have accurately recognized the game and its parameters. The available evidence suggests that the airplane crashed into an empty field, rather than a populated target, because the passengers responded aggressively, challenging the terrorists for control of the airplane, and — although not successful in regaining control — preventing them from accomplishing their nefarious goal.

These passengers, as is well known, had a crucial information advantage that the passengers on the New York airplanes could not possibly have had. Their airplane was in the air at the time of the World Trade Center attacks and after it was hijacked some passengers were able to contact trusted individuals on the ground by cellular telephone. From them, they learned of the other airplanes' fates. Such highly relevant information would contradict a bias towards seeing the game as standard hawk-dove, and consequently the passengers would have analyzed whatever signals the terrorists gave them — presumably the same signals given the victims on the other airplanes — in a very different context, one in which the signals appeared, accurately, to be manipulative rather than informative. Recognizing that the game the hijackers had actually offered was one in which the return for playing dove was exceptionally costly, the passengers recognized that fighting back was the dominating response to the terrorists' aggression. If they played dove, they and many others would die. But if they responded aggressively there was at least *some* chance of saving themselves and an even better chance of saving many potential victims on the ground.

To say that a strategy is rational is not, of course, to diminish the heroism of those passengers who made this choice. They not only *recognized* the best choice that was available to them in ghastly circumstances, but they also *acted* on it, rather than waiting and hoping that someone else would act for them. Fighting back allowed these passengers to die fighting, to frustrate the aims of their enemies, to save the lives of hundreds or even thousands of people, and to be remembered unambiguously as heroes. But they were able to recognize the game's parameters *only* because the tragedy of the first two crashes, relayed to them by reliable sources, provided critical information about the terrorists' intentions.

The cost of misinformation in deadly conflicts

We have less information available about the situation of the passengers on the airplane that was crashed into the Pentagon. It was also in the air at the time of the World Trade Center attacks, and it is known that at least one of the passengers, Barbara Olson, used a cellular telephone to inform her husband, the United States Solicitor General, Theodore Olson, that they had been hijacked, but we do not know what information he gave (or was able to give) her. What is known is that the passengers did not manage to prevent the terrorists from achieving their goal. We see four possibilities: they may not, despite the phone call, have received the information they needed to recognize the game accurately; they may not have had time to organize an aggressive response; they may have responded aggressively but failed to regain or to disrupt control of the airplane; or they may have been in the act of responding aggressively when the airplane hit the Pentagon.

By this analysis, passengers on all of these flights responded rationally, given the information that was available to them. But only those on the flight that crashed in Pennsylvania clearly had the necessary information to recognize the game the hijackers were *actually* playing, as well as the time that was necessary to act on that information. The answer to our second question, then, is that the consequences of success or failure in recognizing the game were enormous — no less than the difference between life and death for thousands of people. If this analysis is correct, what can we learn from the outcomes of these events?

The immediate policy response by the Federal Aviation Administration to the attacks was to require all airplanes that were not already in the air to remain on the ground, and to direct all of the roughly five thousand airborne flights to land at the nearest airfield, with all international flights diverted to Canada. This response was based on the reasonable, and probably accurate, assumption that more hijackings were intended. But what information was provided to passengers and crews of those in-flight airplanes so vulnerable to attack? We do not know exactly what pilots were told, but anecdotal accounts suggest to us that passengers were told only that their airplane was going to have to land immediately, with no information being passed on about the circumstances that made such action necessary.

Perhaps there was concern about causing panic among the passengers, but our analysis suggests that a better policy might have been to tell them *exactly* what was going on. The suicide hijackings that succeeded did so primarily because of the passengers' lack of information, and the policy response did nothing to correct that dangerous condition. Telling everyone just what the situation was would have provided crucial information not only to innocent passengers, but also to any terrorists who *were* on board and preparing to act. That information would have had the potential dramatically to lower the terrorists' estimate of the probability of achieving their goal — taking over the airplane and crashing it into a high-value target — by making clear to them that the passengers could now correctly assess the game, and thus would be more likely to fight back against anyone who stood up and acted aggressively.

Although directing each flight to land immediately at the nearest airport without providing any explanation would also have changed the terrorists' expected-value calculation — by reducing the possible V they could obtain — it would not necessarily lower it enough to deter them from acting. If a “most desirable” target became unattainable, an alternative could have been better than nothing; there is speculation that even the Pentagon was a second-best choice of terrorists whose intended target may have been the White House — and any city with an airport large enough to land one of the diverted jets may have offered “good enough” targets. Non-obvious targets can, in fact, be more desirable than we might at first think. Just as Timothy McVeigh's bombing of a federal government building in Oklahoma City was especially shocking because it targeted a place that no one had expected, so crashing a jumbo jet into a less obvious target than New York City or Washington, D.C., could have made Americans feel there was no place they could be safe — and creating *this* feeling is the classic goal of terrorism.

We cannot say that providing full information to passengers would necessarily have led to more desirable outcomes on September 11. Possibly, innocent actions could have been mistaken by passengers and crew for a terrorist's hawk move, with real harm being done *by* innocent passengers *to* innocent passengers — especially, perhaps, to those perceived as looking “Middle Eastern.” This risk, however, is to be com-

The cost of misinformation in deadly conflicts

pared with the gains that might be realized by dissuading hijackers from making a move, given anticipated resistance.

As we noted, the Federal Aviation Administration's decision to ground all flights was based on the assumption that more hijackings were intended for that day. It seems likely that there was at least one cadre of hijackers who never got off the ground. When it was announced that no more flights would be allowed to take off, several passengers on an airplane at John F. Kennedy International Airport reportedly huddled together for a heated discussion, then rushed off the airplane as quickly as possible when it returned to the terminal. Assuming these men were, in fact, terrorists, the grounding of their flight likely prevented them from killing more innocent people. To them, in a position to know the real meaning of the day's events, the grounding of their flight likely provided information not available to the other passengers — namely, that at least some people knew the true nature of their game. Despite reports that the terrorists, when taking flying lessons, had been more interested in learning how to control an airplane in flight than in learning to land or take off, it is not impossible that they might have been able to get the airplane off the ground if they had tried or that they could have forced the pilots to do so — if and *only if* those pilots also lacked full information about the events of that morning. Unfortunately, we lack knowledge about why they did not try to take off, but if they were terrorists with the ability to get aloft, their decision-making certainly would have incorporated the new information that the grounding imparted to them. Their situation, however, was much different than that of any hijackers who might have been in the air. Knowing that authorities were now forewarned about suicidal hijackings would have been much more of a deterrent to terrorists who still faced the difficult task of taking off than to those with the comparatively simple task of just crashing the airplane.

The only hijacked flight in which passengers thwarted the terrorists appears to have been the only one in which they had accurate information about the game that had been forced upon them. Fortunately, none of the airborne airplanes ordered to land was hijacked and crashed. While we recognize that aviation authorities had to make instantaneous decisions under

great stress while we have the luxury of time to analyze the events, we think this result may have been due more to luck or the inability of the terrorists to adapt to unforeseen circumstances. We conclude that the best strategy would have been to instruct crews on as yet *unhijacked* flights to give their passengers full information. This would not have guaranteed preemption of an attack or even such an attack's failure. But increasing the probability that passengers and crew would respond to hawk moves with individual or collective hawk moves of their own — and making any terrorist aware of that increased probability — *could* have changed expected-value calculation sufficiently to keep suicide terrorists in their seats, albeit with the expectation of flying, fighting, and dying another day.

References

1. Norton Long, "The Local Community as an Ecology of Games," *American Sociological Review*, 1958, 66:251-261.
2. Esther Goody, "Social Intelligence and Interaction" (Cambridge: Cambridge University Press, 1995).
3. Andrew Whiten and Richard Byrne, "The Machiavellian Intelligence Hypothesis" in *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*, R. W. Byrne and A. Whiten, editors (Oxford: Clarendon Press, 1988), pp. 1-9.
4. John Orbell, Tomonori Morikawa, and Nicholas Allen, "The Evolution of Political Intelligence: Simulation Results," forthcoming in *British Journal of Political Science*.
5. Gerd Gigerenzer, "Domain-Specific Reasoning: Social Contracts, Cheating and Perspective Change." *Cognition*, 2000, 43:127-171.
6. Matthew L. Wald, *New York Times*, October 16, 2001, p. B9.
7. John R. Krebs and Richard Dawkins, "Animal Signals: Mind-Reading and Manipulation?" in *Behavioural Ecology: An Evolutionary Approach*, J. R. Krebs and N. B. Davies, editors (Oxford and Malden MA: Blackwell Science, 1984).
8. John Maynard Smith and G. R. Price, "The Logic of Animal Conflict." *Nature*, 1973, 246:15-18.
9. Tomonori Morikawa, James E. Hanley, and John Orbell. "Cognitive Requirements for Hawk-Dove Games: A Functional Analysis for Evolutionary Design," *Politics and the Life Sciences*, March 2002, 21:1, pp. 2-10.
10. W. D. Hamilton, "The Genetical Evolution of Social Behaviour, I," *Journal of Theoretical Biology*, 1964, 7:1-16
11. Lelyveld, Joseph, "All Suicide Bombers Are Not Alike," *New York Times Magazine*, October 28, 2001.