

On the Economics of Anonymity

Alessandro Acquisti¹, Roger Dingledine², and Paul Syverson³

¹ SIMS, UC Berkeley (acquisti@si ms. berke ley. edu)

² The Free Haven Project (arma@mi t. edu)

³ Naval Research Lab (syverson@i td. nrl . navy. mi l)

Abstract.

model that incorporates many of them. In Section 4 we give some simplifying assumptions and draw conclusions about certain scenarios. Sections 5 and 6 describe some alternate approaches to incentives, and problems we encounter in designing and deploying strong anonymity systems.

2 The Economics of Anonymity

Single-hop web proxies like the Anonymizer protect end users from simple threats like profile-creating websites. On the other hand, users of such commercial proxies are forced to trust them to protect tra c information. Many users, particu-

1. Costs of sending messages through the anonymous system, c

where (\cdot) , (\cdot) , and (\cdot) are unspecified functional forms. The payoff function $u(\cdot)$.

the message). Thus initially we do not consider the strategy of choosing to be a bad node, or additional honest strategies like creating and receiving dummy traffic.

We represent the game as a simultaneous-move, repeated game because of

nodes. We also assume that all agents perceive the same level of anonymity in the

the following way:⁸

Agent i / Agent j	a_j^h	a_j^s	a_j^n
a_i^h	A_i, A_j	D_i, B_j	E_i, C_j
a_i^s	B_i, D_j	F_i, F_j	G_i, C_j
a_i^n	C_i, E_j	C_i, G_j	C_i, C_j

a model where each player plays a large set of identical players, each of which is "infinitesimal", i.e. its actions cannot affect the payoff of the first player. We define the payoff of each player as the average of his payoffs against the distribution of strategies played by the continuum of the other players. In other words, for each agent, we will have: $u_i =$

can be also compared to [15], where the paradox of informationally efficient markets is described.¹¹

The problems start if we consider now a different situation. Rather than having a continuous distribution of valuations v_{a_i} , we consider two types of agents: the agent with a high valuation, $v_{a_i} = v_H$, and the agent with a low valuation, $v_{a_i} = v_L$. We assume that the v_L agents will simply participate sending traffic if the system is cheap enough for them to use (but see Section 6.3), and we also assume this will not pose any problem to the v_H

2. *"Special" agents*

of the main open problems in the design of any decentralized anonymity service. The Advogato trust metric [16] and similar techniques rely on humans to

6.3 Bootstrapping The System And Perceived Costs

Our models so far have considered the strategic choices of agents facing an already existing mix-net. We might even imagine that the system does not yet exist but that, before the first period of the repeated-game, all the players can somehow know each other and coordinate to start with one of the cooperative equilibria discussed above.

But this does not sound like a realistic scenario. Hence we must discuss how a mix-net system with distributed trust can come to be. We face a paradox here: agents with high privacy sensitivity want lots of traffic in order to feel secure using the system. They need many participants with lower privacy sensitivities using the system first. The problem lies in the fact that there is no reason to believe the lower sensitivity types are more likely to be early adopters. In addition, their *perceived* costs of using the system might be higher than the real costs¹²

9. Whitfield Diffie and Susan Landau. *Privacy On the Line: The Politics of Wire-tapping and Encryption*. MIT Press, 1998.
10. Roger Dingledine, Michael J. Freedman, David Hopwood, and David Molnar. A