WP. 40
ENGLISH ONLY

UNITED NATIONS STATISTICAL COMMISSION and ECONOMIC COMMISSION FOR EUROPE
CONFERENCE OF EUROPEAN STATISTICIANS

EUROPEAN COMMISSION
STATISTICAL OFFICE OF THE EUROPEAN COMMUNITIES (EUROSTAT)

Joint UNECE/Eurostat work session on statistical data confidentiality
(Bilbao, Spain, 2-4 December 2009)

Topic (vii): Risk/benefit analysis and new directions for statistical disclosure limitation

AN APPLICATION OF GAME THEORY TO UNDERSTANDING STATISTICAL DISCLOSURE EVENTS

Invited Paper

Prepared by E.C. Mackey and M.J. Elliot (University of Manchester), United Kingdom

An Application of Game Theory to Understanding Statistical Disclosure Events

Mackey, E. C* and Elliot, M. J.* [1]

*Centre for Census and Survey Research University of Manchester, UK.

Abstract: To fully account for the risk of a statistical disclosure occurring requires that we develop a better understanding of how a disclosure might occur and what its consequences might be. To do this we need to consider not just the target data but also the environment in which that data is produced and released. Through this we can identify and explore the events leading up to and following from a (claim of) disclosure. That is we can move beyond measuring disclosure risk to conceptualising and systematically representing disclosure events in their entirety. In this paper we show how it is possible to apply a game theoretic reasoning to model disclosure events to examine how key agents involved in creating a disclosure might interact to bring about particular outcomes. The paper gives a brief introduction to game theory and to how it might in general be applied to the disclosure event situation and illustrates this with some examples.

1 Introduction

In Statistical Disclosure Control one of the main concerns has been: how to restrict the data for release to ensure a balance is reached between confidentiality and data utility (for various perspectives on this issue see: Duncan et al 2001; Kennikel and Lane 2006; Purdam and Elliot 2007; Shlomo 2007; Trottini 2003). To make decisions about how best to restrict data, one first assesses the risk of a disclosure occurring. This requires, as a starting point, that we establish an understanding of the events that might lead to a disclosure such as an intrusion attack or a spontaneous disclosure[2]. There are three things that we need to understand about disclosure events: 'how they might actually happen', 'what they are', and 'what might their consequences be'? The SDC literature has focused heavily on the second of these and there is a pressing need to strengthen our understanding of the other two components. However, ascertaining how a disclosure event might arise and further play out is far from straightforward. This is because data is released into a complex environment of which we have only partial knowledge that is constrained by the predominate focus on the data to be released rather than the environment itself. In order to gain an insight into the data environment we have to shift our focus away from the data towards human agency, because it is only when we look at the actions

[1] The research reported here was carried out supported by funding from the Economic and Social Research council and the Office of National statistics under CASE Studentship award S422134054.

[2] A spontaneous disclosure refers to a situation whereby someone looking at the visible information in a dataset is able to learn new information. It may come about from the release of a table cell count of zero (see Smith and Elliot 2008 for example).

of, and interactions between, the key agents in the dissemination settings that we can begin to understand how disclosure events are created and their outcomes shaped.

2 Disclosure Events

Let us begin by saying what we mean by a disclosure event. It can be thought of as a core occurrence arising from an interaction between agents[3] such as a data stewardship organisation (DSO)[4] and intruder that leads to a disclosure claim. The disclosure claim may or may not be based on an actual disclosure and the claim may or may not be public.

The point at which action commences, with the DSO making decisions about a potential data release, to when an agent believes he can make a disclosure claim can be described as a disclosure event scenario. The use of scenarios as a tool for developing a better understanding of how a disclosure might occur is not new in SDC. The work of Paass (1988) Mokken et al (1992) Lambert (1993) and Elliot and Dale (1999) established scenario analysis as a key part of disclosure risk assessment. The usefulness of scenarios lies in their potential to illuminate a range of conditions required to bring about a potential disclosure claim. However, until now they have been used in a uni-dimensional and static way and though this has led to a practically useful analysis of the actions of a single agent (referred to here as an intruder) it has not crucially allowed for an examination of the inter-action that occurs amongst key agents when, for example, a DSO releases data and a data user or would be intruder accesses it. An analysis of the potential interactions between agents is fundamental to our understanding of how a disclosure might actually arise and play out. This is because the actions of agents such as the DSO and data intruder are interdependent meaning that each agent's actions affect the actions of one or more other agents. Thus, to understand the probable actions of an intruder we must consider them in relation to the actions of the DSO and vice versa. We might therefore hypothesise that the best actions of the DSO are based on its perceptions of the best actions of the intruder. Similarly, the best actions of the intruder are based on his perceptions (and partial knowledge) of the best actions of the DSO. In practice, this means that the DSO directs all its efforts to avoid a disclosure by attempting to thwart the efforts of the intruder whilst the intruder acts to direct all his efforts towards counteracting those of the DSO in an effort to bring about a plausible disclosure or plausible

[3] Agents in the data dissemination setting include also data users, the general public, media, specialist interest groups and SDC specialists. Although these agents might at first glance appear to be on the periphery, they have a key role in helping us understanding how a disclosure event might play out. These agents might be drawn into a direct interaction with the DSO as a primary player or indirectly as a secondary player.

[4] After Duncan et al (2010) we use data stewardship organisation to refer to any organisation that has responsibilities for data dissemination processes.

disclosure claim. In short, the interdependent nature of agents' actions requires that we focus on the interaction between agents, rather than the actions of any single agent, if we are to fully comprehend the disclosure risk problem. Not least because by exploring disclosure event scenarios of multi-agent action and inter-action we have the potential to: (i) more fully delineate the range of conditions required to bring about a disclosure claim and (ii) identify who the key agents might be and how they might interact to bring about particular outcomes.

We can identify three types of disclosure event for which one could develop event scenarios: *de facto* intrusion, spontaneous recognition and spontaneous disclosure. We shall refer only to the disclosure event of intrusion in the remainder of the paper.

The details or specifics of a disclosure event scenario will be shaped by the DSO's data release decisions, i.e. on the type of data released, SDC methodology, level of geographical detail and the provision of metadata, and the other key agents' counter-responses to these decisions. For example, a DSO's releases a dataset (the make up of which it determines) that is then accessed by a would-be intruder. The intruder has, we would argue, two possible response frameworks that he could employ. Which, if any, of these response frameworks he could successfully utilise to bring about a statistical disclosure (or make a disclosure claim) will be determined[5] not just by the DSO's data dissemination policies but also the intruder's own skills, knowledge, access to external data sources and motivations. The intruder's response options are: directly exploit the released data using other related sources of data (public or restricted access); manipulate public perceptions by exploiting the released data[6], threaten the DSO or some other agent.

So far we have considered how we might describe our agents' interactions. However, if we wish to move beyond a descriptive examination to consider how agents' interactions might play out, and what the consequences of particular outcomes might be, we need a framework for modelling one agent's action in relation to another agent's action. One such framework is Game Theory.

[5] There are other factors, external to the released dataset, that will also influence which, if any, tactic an intruder could successful employ to bring about a disclosure claim, these are related to the data environment i.e. an intruder's ability to exploit useful external data sources and/or manipulate public perceptions.

[6] In order to do this an intruder would need to make others (i.e. General Public) believe a disclosure had occurred regardless of whether the claim can be substantiated.

3 Game theory[7]

'Whenever the choices by two or more distinct decision-makers have an effect on each others' outcomes, the interaction between them is game-theoretic in nature', (Dickson and Nowaczyk 2004:1).

Given the interdependent nature of the relationship between agents in the data dissemination setting, a game theoretic framework, we suggest, could be applied to help us understand and analyse the interactions between our agents. More precisely, it can help us: (i) develop plausible disclosure event scenarios through which to explore the likely outcomes and consequences of agents' interactions and; (ii) with the use of real data assess the likelihood (high, medium, low) of a particular disclosure event occurring.

As Levine (2000) says game theory can be described as a 'science of strategy'. It uses mathematics to determine the actions an agent[8] should or will take, in an interaction, to obtain the best outcome for him whilst taking account of the actions of other agents.[9] Within game theory each agent has a strategy profile i.e. a list of all his possible actions (strategy choices). Each agent's strategy choices in a game, of which there may be many, are commonly represented by numerical payoffs based on their preferences. These preferences are ordered and mapped onto real numbers, via a utility function. For example, suppose agent one prefers A to B and B to C, the ordinal[10] utility function maps the highest ranked preference with the largest number and then the next highest ranked preference with the next largest number and so on. This may be represented like this: A >> 3, B >> 2, C>>1. The aim of each player is to end the game with the largest payoff. This, however, will only be partially within each agent's control since the outcome, and therefore payoffs, will depend upon the strategy choices of all the agents. Agents' behaviour is defined as strategic because they take into account this interdependence when they choose which of their strategies to play.

[7] In producing this summary of game theory we have drawn on several general texts on the topic. Specifically; Dixit And Nalebuff 1991; Binmore 2007; Straffin 1993; Camerer 2001; 2003; Wright 2002; McCain 1997.

[8] An agent or player can be virtually any entity: a person, a nation or business (Camerer 2001).

[9] Game theory does not attempt to explain why agents act in the way that they do although this does not mean that game theorists are not concerned with analytical type questions. They are, it is just that these types of questions are directed at the game level. For example, one might ask why certain game conditions produce particular outcomes.

[10] An ordinal utility function tells us about the ranking of a player's preference, it does not inform us about the strength of his preference. To account for an agent's strength of preference we would require a cardinal utility function. However, in general, it is questionable whether one can apply a cardinal utility to non-economic outcomes (see Fang et al 1993 for discussion).

In general, it is also assumed that players are rational and therefore play to maximize their utility. The notion of rationality is a core concept in game theory as it underpins both descriptive and predictive analysis about how a game should/will be played. It is thought that if one can work out the best response of each player to the actions of his opponent then one can find the game's (unique) solution[11]. However, rational play involves an individual making complicated decisions about: (i) how to choose a strategy with a desirable outcome for oneself whilst knowing that one's opponent is also trying to choose a desirable outcome and (ii) decisions about whether one's interests coincide or conflict with one's opponent. Not unsurprisingly, in practice, players do not always act to maximize their utility. As Straffin points out: 'In the real world, it is doubtful that all players will play rationally', (1993:4).

The idea that players don't always play rationally raises something of a problem. This is because, if the game's outcome is dependent on the choices each player makes, the presence of a player that does not play to ensure his preference is never violated changes what a rational opponent should do. There have been attempts to take account of a less stringent model of rationality but so far it has had limited success. Thus, a model of economic rationality continues to be the most widely applied model in traditional game theory. Commonly, a weaker assumption of rationality is used: that a player is motivated to maximize their payoff and it is this assumption that we will employ here.

4 A framework for analyzing disclosure event games

The branch of game theory that provides the framework for the disclosure event game described later is non-cooperative or traditional game theory. Contrary to what one might imagine from its title it is concerned with interactions where there is a basis for cooperation as well as interactions where there is not. The crucial element is that when agents do decide to co-operate this cooperation is self-enforcing. There are two types of games that will be of interest to us: (i) coordination games where players share a mutual interest in playing the same strategy and (ii) anti-coordination games where players share a mutual interest in playing different strategies.

Let us consider the DSO-intruder interaction in more detail. In this interaction, we hypothesise that, the DSO makes the first move by setting the initial conditions of the game and releasing data into the public domain. The intruder can respond by making two consecutive moves each a dichotomy of decision options. So in move 1, the intruder can either attempt to identify and/or disclose new information about

[11] Solution concepts can be thought of as rules that predict the actions of players. There are several solution concepts such as Von Neumann and Morgenstern 'Minimax Theorem' for zero sum games and for non-constant sum games John Nash's Nash Equilibrium.

respondents or attempt to make others believe he has identified and /or disclose new information about respondents. In move 2, he can either publish his disclosure (attempt/success) or seek no publicity for his actions. The DSO can also make a subsequent move, it could either: make a second move at the beginning of the game to pre-empt the actions and/or the consequences of a intruder's actions by talking proactively about the issue of disclosure or make a fourth move (later in the game) by talking/not talking about disclosure in reaction to the intruder's move of publicizing a disclosure. This series of moves is illustrated in figure 1.

Fig 1: DSO-intruder interaction

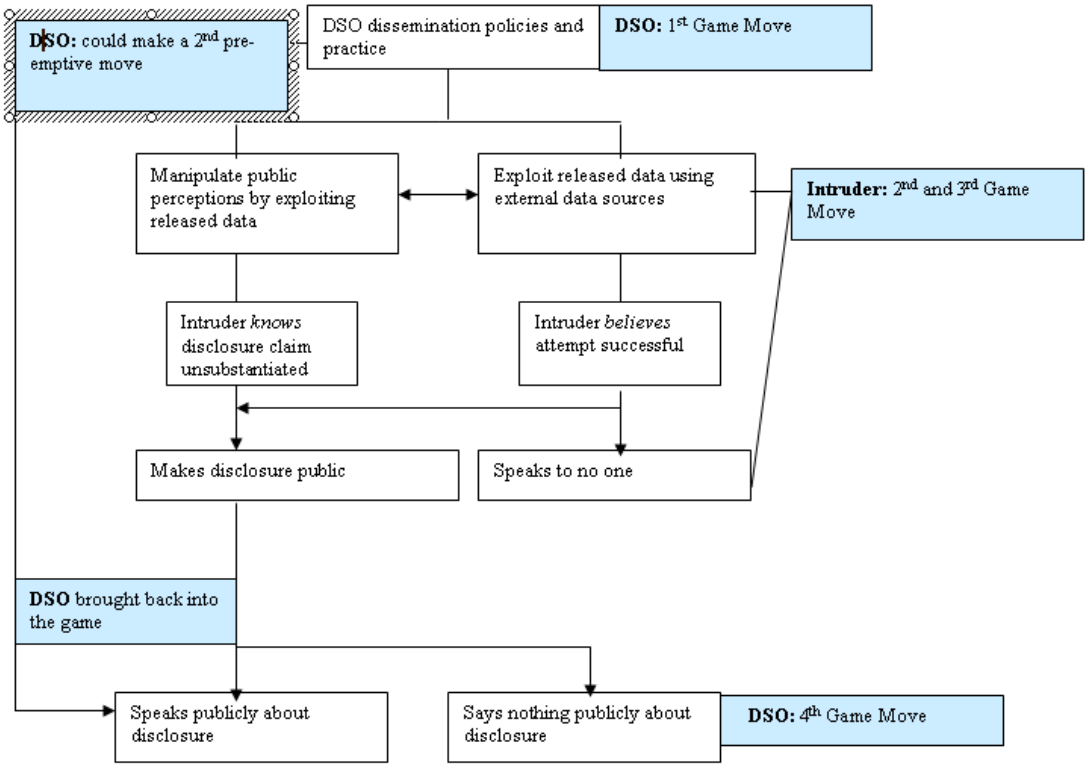

From this basic interaction, between the DSO-Intruder, we shall set out one game (many more will be possible) to explore the possible outcomes of a combination of moves by our agents. The game outlined here is an exemplar of possible interaction. We shall model a simple two-player game applying a simplifying assumption to each of the agents to treat them as representative agents of all similar agents.

All games start with a condition of play. This is established in the data dissemination processes that the DSO utilises when releasing a dataset. The intruder makes the next move aware of the move that has gone before. In figure 1, the intruder's 2nd move involves him making decisions about how best to respond to the DSO efforts to maintain data confidentiality. To examine this combination of moves and assess the

likelihood of the intruder's success requires the use of real data. For this exemplar, we shall bypass modelling this move and model the next on the assumption that the intruder has been successful in his endeavour to bring about a disclosure. The intruder as we have said has two options, either to 'publicize his disclosure/claim' or 'avoid publicity': so what happens next?

Before we can go any further with our game we must first clarify how we will determine our agent's strategy options and their preference for one option over another. This has two components: (i) the identification of possible strategies an agent might have in a particular interaction and (ii) the determination of which of the available strategy options an agent might prefer. In terms of identifying possible strategies we ask 'what does the agent hope to achieve' by his actions. We shall illustrate this once we set out our game later. In terms of determining an agent's preference for a strategy we look at approaches, or roles an agent might take in an interaction, which would lead him to favour one action over another. We shall consider here just two for each agent. These approaches do not amount to an exhaustive list of all possible approaches that an agent might take rather they serve as an illustration of how a single interaction can lead to multiple possible games and associated outcomes. For an intruder, we consider the following approaches: *covert* and *overt*. We shall describe the intruder's approach as covert if he wants to keep his actions hidden and overt if he wants to make his actions known. For the DSO, we consider the following approaches: *engaging* and *reserved.* I shall describe the DSO approach as reserved if its objective is to avoid been drawn into a public discussion on disclosure and engaging if it is willing to take part in a public discussion on disclosure.

By considering multiple approaches for each agent we can explore multiple variations of the one interaction. For the approaches outlined previously we have four potential games:
(**1**) The intruder takes a covert approach whilst the DSO approach is reserved.
(**2**) The intruder takes an overt approach whilst the DSO approach is reserved.
(**3**) The intruder takes an overt approach whilst the DSO approach is engaging.
(**4**) The intruder takes a covert approach whilst the DSO approach is engaging.

We shall consider here only one variation of the DSO-Intruder interaction, game (2), when the intruder takes an overt approach and the DSO a reserved position. This variation of the DSO-intruder interaction as we shall now illustrate will lead to further games and new interactions. So how might we understand this combination of preferred action? A DSO can still respond to the intruder's action without formally talking publically, for example, by pursuing a prosecution. Let us model this interaction.

The DSO has to weigh up the benefit of pursuing a prosecution, for example that it sends a clear message to anyone thinking about breaching data confidentiality, against any costs such as an increase in negative publicity. We could hypothesise that the DSO will be motivated to: (i) deter and prevent future breaches in data confidentiality and (ii) limit the damage to its reputation. It has two possible strategy choices: (1) Threaten a prosecution; (2) Pursue a prosecution. In response to the DSO move the intruder, we might hypothesise, will be motivated to either (i) avoid prosecution and/or (ii) avoid punishment. For current purposes we will consider that he also has two options available: (1) Withdraw the disclosure claim from the public domain (2) Defend disclosure by making a public interest claim. In order to determine the players' preference for one strategy over another let us consider the approaches they may take. For the DSO we shall consider: *damage limitation* and *prevention*. We shall describe a DSO's approach as damage limitation if its objective is to limit the harm to its reputation and prevention if its objective is to ensure future would be intruders are discouraged. For the intruder we shall consider the approaches: *offensive* and *defensive*. We shall describe an intruder's approach as offensive if his objective is to pursue his original position regardless of any threat made by the DSO and defensive if he retreats from his original position in the face of a threat of prosecution. As before there are four possible variations to this game:
(**1**) The DSO takes a damage limitation approach, the intruder an offensive approach.
(**2**) The DSO takes a prevention approach, the intruder an offensive approach.
(**3**) The DSO takes a damage limitation approach, the intruder a defensive approach.
(**4**) The DSO takes a prevention approach, the intruder a defensive approach.

To play the game we need to first establish our agents' preferences for their strategy options, we may write them thus. If the DSO takes a damage limitation approach it will prefer to threaten prosecution to actually pursuing it and its utility function can be written as: Threaten prosecution >>2; Pursue prosecution >>1. Alternatively, if the DSO takes a prevention approach it will prefer the reverse and its utility function can be written as: Pursue a prosecution >>2; Threaten a prosecution >>1.

If the intruder takes an offensive position he will prefer to defend his claim to withdrawing it and his utility function will be: Defend claim >>2. [DC]; Withdraw claim >>1. [WC]. Alternatively, if the intruder takes a defensive approach then he will prefer the reverse and his utility function can be written as: Withdraw claim >>2. [WC]; Defend claim >>1. [DC]

We can model the DSO-intruder interaction using a game tree graph (see figure 2) as we have an order of play: (i.e. the intruder publicizes his disclosure and the DSO responds). The different variations of this interaction are labelled games 1 to 4 accordingly. The first number in the boxes belongs to the DSO and the second number belongs to the intruder.

Fig 2: DSO-Intruder interaction
KEY: 2 represents a player's most preferred option; **1** represents a player's least preferred option.

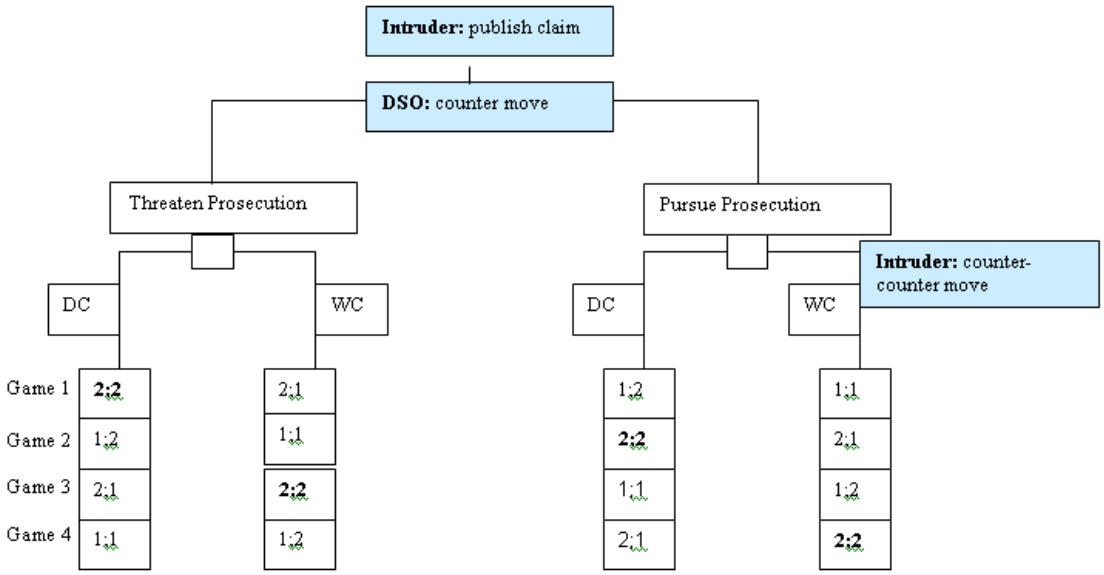

To determine the most likely outcome of each of these games we need a solution concept to help short circuit the agents reasoning: 'the DSO thinks that the intruder thinks that the DSO thinks that……'. One such solution concepts is Nash Equilibrium (NE). It is one of the most commonly used concepts based on the idea of equilibrium in the physical environment (Crawford 2003). An NE is said to occur where there is a set of strategies with the property that no player can benefit by changing his strategy unilaterally. Many games have more than one NE. This situation is regarded as problematic because theorists would like to be able to say with some precision which outcomes in a game should/will occur. In cases where there is more than one NE another process may be employed that will enable players to converge on the same equilibrium, when they cannot communicate. One such process is Schelling's (1960) theory of focal points, which suggests that players may coordinate on a particular equilibrium by using information that is abstracted from outside the game as well as information from inside of it.

A game's outcome will be dependent upon the approaches taken by each of the players. For example, in game 1 when the DSO prefers to limit the damage to its reputation and the intruder prefers an offensive approach the agents best payoff occurs when the agents play opposing strategies i.e. DSO threatens prosecution but the intruder maintains his position. The DSO requires the intruder to take his least prefer option i.e. 'to change his position' to obtain its best outcome, as illustrated in

Game 1. [2:1]. The intruder however has the advantage of being able to see the DSO move before he makes his own. So why should he take his least preferred strategy? In this game scenario the balance of power in the game lies with the intruder. This is related not just to the agents' order of play but also the approach taken by the DSO. Because threatening rather than pursuing a prosecution the DSO strategy may not be enough of a sanction to influence the intruder choice of strategy. To force the intruder to reconsider his approach the DSO needs to reconsider its approach also. By pursuing a prosecution the DSO has the potential to influence both the intruder's actions (see game 4) and those of future would be intruders. The DSO however will not in all likelihood be able to dissuade all types of intruder from defending a claim (in particular, those who are specifically motivated to generate as much publicity as possible for their actions). This outcome might be particularly problematic for the DSO if it continues to take a reserve approach and not talk publicly about the intruder's disclosure/claim. In all likelihood it would initiate a new game this time with another agent: 'the general public'.

5 Concluding remarks

In this paper we have outlined the case for expanding the scope of disclosure risk assessment beyond the mechanics of disclosure to the environment in which disclosure events might take place. We have shown how game theoretic reasoning might shed some light on the complexity of this environment and have provided a simple example of what a disclosure event game might look like.

Another advantage of the approach outlined here is that it has wider application to data types than the official microdata and tabular aggregates that are the focus of orthodox disclosure risk analyses. Confidentiality issues surrounding social network data, administrative datasets and qualitative data would all be amenable to treatment using game theoretic techniques. The possibility of encapsulating a larger range of data types and data confidentiality issues in a single general framework is very attractive.

Clearly, real disclosure event scenarios will be far more complex than space in this paper has allowed to us to explore and the interested reader is directed to Mackey (2009) for further discussion of some of the issues. However, we feel that the approach has something significant to offer the field and will pursuing it further in future work.

References

Binmore, K. (2007). Game Theory: A Very Short Introduction. Oxford University Press.

Binmore, K. (1994). Game Theory and Social Contract: Playing Fair. MIT Press.
Camerer, C. F. (2003). *Behavioural Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
Camerer, C. F. (2003). Behavioural Studies of strategic thinking in games. *TRENDS in Cognitive Sciences*. Vol.7. No.5. May 2003. Elsevier. pp225-235.
Camerer, C. F.; Teck-Hua Ho.; Chong, J.K. (2001). *Behavioural Game Theory: Thinking, Learning and Teaching*. http://www
Crawford, V. (1997). Theory and Experiments in the Analysis o Strategic Interactions in Kreps, D. and Wallis, K. (eds). *Advances in Economics and Econometrics: Theory and Applications.* Seventh World Congress. Vol1. Econometrics Society Monographs No 27. Cambridge, New York. Cambridge University Press. Pp206-242.
Dickson, E. and Nowaczyk, P. (2004). *Introduction to Game Theory in Political Science*. Online at http://www.nyu.edu/gsas/dept/politics/undergrad/syllabi/v53.0840-dickson-/04.pdf.
Dixit, A. and Nalebuf, B. (2002). *Game Theory*. The library of Economics and Liberty. The Concise Encyclopaedia of Economics. Online at *http://www.econlib.org/library/Enc/Gametheory.html*.
Duncan, G. T., Keller-McNulty, S. A., and Stokes, S. L. (2001): "Disclosure Risk vs. Data Utility: The R-U Confidentiality Map", Los Alamos National Laboratory Technical Report LA-UR-01-6428.
Elliot, M.J. and Dale, A. (1999). Scenarios of Attack: The Data Intruder's Perspective on Statistical Disclosure Risk. *Netherlands Official Statistics*. Special edition, Spring, pp 6-10.
Elliot, M.J. (1996). *Attacks on census confidentiality using the Sample of Anonymised Records*. Presented at the Third International Seminar on Statistical Confidentiality, Bled, Slovenia, October.
Fiori, S. (2005). Simon's Bounded Rationality. Origins and use in Economic Theory. Working Paper No. 09/2005
Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press.
Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press.
Kennickell, A. B. and Lane, J. (2006): "Measuring the impact of data protection techniques on data utility: Evidence from the Survey of Consumer Finances", in J. Domingo-Ferrer, ed. *Privacy in Statistical Databases* (Lecture Notes in Computer Science):291-303.New York: Springer-Verlag.
Lambert, D. (1993). Measures of Disclosure Risk and Harm. *Journal of Official Statistics*. Vol. 9, No. 2, pp 313-331.
Levine, D. (2000). *An overview of Game Theory*. Online at http://www.levine.sscnet.ucla.edu/general.htm, , accessed 9.02.2007.
Marsh, C., Skinner, C., Arber, S., Penhale, B., Openshaw, S., Hobcraft, J., Lievesley, D. and Walford, N. (1991). The Case for Samples of Anonymised Records

from the 1991 Census. *Journal of the Royal Statistical Society* 154(2) pp 305-340.
McCain, R. (1997). *Strategy and Conflict: An introductory Sketch of Game Theory*. Online at http://william-king.www.drexel.edu/top/eco/game/game.html, accessed 10/01/20008.
Mokken, R. J., Kooiman, W. J., Pannekoek, J. and Willenborg. (1992). Disclosure Risk for Microdata. *Statitica Neerlandica*, Vol.46 pp49-67.
Paass, G. (1988). Disclosure Risk and Disclosure Avoidance for Microdata. *Journal of Business and Economic Studies*. Vol. 6, No. 4, pp 487-500.
Purdam, K. and Elliot, M. J. (2007): "A Case Study of the Impact of Statistical Disclosure Control on Data Quality in the Individual UK Samples of Anonymised Records", *Environment and Planning A*. 39, 1101-1118
Rasmusen, E. (200). An Introduction to Game Theory. Blackwell Publishing
Rubinstein, A. (2005). Behavioural Economics. Advances in Economic and Econometrics Theory and Applications, Ninth World Congress.
Shlomo, N. (2007). *Statistical Disclosure Control Methods for Census Frequency Tables*. *International Statsitics Review, Vol. 75, Number 2*, 199-217.
Smith, D. and Elliot, E. (2007). *A Measure of Disclosure Risk for Aggregate Data.* Joint UNECE/EUROSTAT work session on statistical data confidentiality. Manchester. UK, 17-19 December 2007.
Straffin, P. (1993) *Game Theory and Strategy*. MMA Mathematics Association of America.
Trottini, M. (2003): "Decision Models for Data Disclosure Limitation", Ph.D. thesis, Carnegie Mellon University, Department of Statistics.
Wright, R. (2002). *Both Sides Now.* The Earthling: Science, Evolution and Politics Explained. *http://www.slate.com*