

Issues in measuring the efficacy of a suspicious activity reports (SARs) regime

by Matthew H Fleming

Measuring the impact of suspicious activity reports (SARs) on crime is, unfortunately, not a straightforward endeavour. Several authors have noted this to be the case, and the present paper seeks to expand on the factors which prevent us from simply “running the numbers.” For the time being at least, a lack of data and some complications to the underlying mechanisms through which SARs affect crime and crime affects SARs present sufficiently high hurdles to the application of econometric methods to estimating a policy impact. While the future may hold more promise (with better data), this suggests that a more pragmatic approach is required. This approach should seek to highlight the value-add which SARs bring to crime detection by leveraging a holistic, networked view of SARs data along with other law enforcement data.

INTRODUCTION

A suspicious activity report (SAR) is a piece of information which alerts law enforcement that certain customer activity—perhaps a series of large out-of-character deposits or the cash purchase of a high-value asset—is in some way suspicious and might indicate money laundering or terrorist financing (and thus criminal and/or terrorist activity). SARs are sent by members of the regulated sector(s) (eg banks, money service bureaux, etc) to national financial intelligence units (FIUs) for processing, and are subsequently (typically) passed to law enforcement for action. SARs regimes in many jurisdictions have undergone considerable change in recent years. For example, money laundering offences have been redefined (typically with the activities comprising money laundering broadened), and regulated sectors have been enlarged to include new businesses, such as lawyers and dealers in high-value goods. Perhaps not surprisingly, and presumably due at least in part to the expansion of regulated sectors, jurisdictions worldwide have witnessed growth in the number of SARs filed year-on-year.

SARs are thought to deter money laundering and predicate offences; facilitate the detection and sanctioning of such crimes after the fact; and disrupt such crimes in progress (see, for example, *Money laundering: review of the reporting system*, KPMG LLP, 2003; Reuter, Peter, and

Truman, Edwin M, *Chasing Dirty Money: the Fight Against Money Laundering*, Washington DC, Institute for International Economics, 2004; Gold, Michael and Levi, Michael, *Money Laundering in the UK: an Appraisal of Suspicion-Based Reporting*, London, The Police Foundation, 1994; and HM Treasury, *Anti-Money Laundering Strategy*, HMT, 2004). Despite the atmosphere of change and the growth in SAR filings, the impact of SARs regimes on crime remains unclear. Indeed, the subject of measuring the efficacy of these regimes on crime has received little overall empirical attention. This paper seeks to review some specific issues surrounding empirical measurement in greater depth, including the question of whether SARs regimes reduce crime through deterrence or facilitate detection after crimes have been perpetrated. The paper begins with a discussion of the challenges surrounding measuring the impact of SARs on crime in general, and then addresses possible metrics for evaluating the specific impact of SARs regimes on crime detection.

MEASURING THE IMPACT OF SARs ON CRIME IN GENERAL

The costs associated with SARs regimes are not trivial (see, for example, KPMG, LLP, 2003, and Reuter and Truman, 2004, above for discussions of cost estimates). So do the benefits of a SARs regime outweigh the costs? The present paper does not seek to compare benefits with

costs. But thinking in benefit/cost terms begs the question: what are the benefits? What specifically should we measure? How should we measure benefits? And, can we, with existing data, even measure the benefits empirically? The likely answer to these questions is no. This section explores these issues in greater depth.

To begin, what should be measured? What is the goal of a SARs regime? Is it to reduce money laundering specifically, which Peter Reuter and Michael Levi call an “intermediate target” (in “Money laundering: a review of current controls and their consequences” in Michael Tonry, ed, *Crime and Justice: an Annual Review of Research*, University of Chicago Press, forthcoming)? Or is it to reduce the incidence of the underlying predicate offences (or might it be to increase asset recoveries)? So the question for empirical analysis then becomes: in determining the efficacy of a SARs regime, should we measure the relationship between SARs and money laundering arrests/convictions, or between SARs and arrests/convictions for predicate offences, or even SARs and asset recoveries? (The present paper does not address the broader anti-money laundering goals of protecting the integrity of the core financial system and reducing terrorism – for more on these topics, see Reuter and Truman, 2004, above). Part of this decision rests on the quality of available data. The challenge here is that while the most direct goal of a SARs regime probably is to reduce money laundering, data on money laundering arrests/convictions are of poor quality (this is discussed in the literature – see Reuter and Levi (2006) and Reuter and Truman (2004) above).

In numerous jurisdictions, financial investigation is a somewhat novel field, with the result that arrests/convictions for money laundering tend to be low at least in part because investigators have not fully understood/embraced the money laundering charge as one worth taking forward. This is changing in the UK, of course, in part due to the efforts of the National Centre for Policing Excellence and Her Majesty’s Inspectorate of Constabulary (see *Payback Time: Joint Review of Asset Recovery Since the Proceeds of Crime Act 2002*, HMIC, 2004 and *Practical Advice on Financial Investigation*, NCPE, 2006). The crime was also often seen by prosecutors as too difficult to prove; legislative changes to the definition of money laundering and the requirements needed to prove the crime have made this less of a barrier, but money laundering charges are often dropped in plea arrangements (see Reuter and Levi, 2006). So while money laundering was clearly practised in the commission of various crimes, this fact is not reflected in historic data on arrests/convictions. In addition, as with all research on crime, data on arrests/convictions represents information on those criminals who were caught, not the broader population of all criminals.

Turning from what should be measured, how might we measure efficacy? To determine the impact of SARs on

money laundering specifically, or crime in general, we would want to run some empirical analyses. Ideally, we would think of crime as a function of SARs and a raft of other variables; gather the data, run the numbers, and then explore whether or not our SARs variable is statistically significant. But, sadly, it is not that easy. In a nutshell, the number of SARs which are filed may affect crime – defined broadly here as money laundering or acquisitive crime – and crime may affect the number of SARs which are filed (of course the number of SARs filed as a variable in empirical analyses should likely be expressed as a rate, perhaps per capita). So the position is:

Changes in the number of SARs filed can affect crime because:

- there may be less underlying crime requiring money laundering because criminals are deterred by SARs;
- there may be less laundering because launderers are deterred by SARs;
- SARs may lead to arrests/convictions, pulling launderers/criminals off the market (and over time, if replacement is less than arrests, we might ultimately see a fall in the number of arrests).

And changes in crime can affect the number of SARs filed because:

- there may be changes in the amount of underlying crime in need of laundering;
- we may see the same amount of crime, but its profitability may change, commanding more/less laundering;
- changes in the legal definition of laundering may imply that more SARs are filed on the same level/kind of behaviour;
- laundering activity may move to unregulated sectors, with fewer SARs thus filed;
- and the organisation of crime may change (behaviour may move from one-off to organised crime), with an unclear impact on the number of SARs filed.

Because of this problem (known as “simultaneity” in econometrics), isolating the crime-reducing efficacy of a SARs regime is far from straightforward. This phenomenon is not new to the study of criminal justice (or of public policy in general). Indeed, the issue has been raised in numerous studies and meta-analyses of criminal justice policies/programs, including examinations of the crime-reduction impact of CCTV, neighbourhood watch, street lighting, police, and prisons (see, for example, Welsh, Brandon C, and David P Farrington, *Effects of Improved Street Lighting on Crime: Protocol for a Systematic Review*, Campbell Collaboration, 2003; Welsh, Brandon C, and Farrington, David P, *Crime Prevention Effects of Closed Circuit Television: a Systematic Review*, Home Office, London, 2002; Levitt, Steven D, “Using electoral cycles in police

hiring to estimate the effect of police on crime”, *The American Economic Review* 87, No 3, (1997) 270-90; Levitt, Steven D, “The effect of prison population size on crime rates: evidence from prison overcrowding litigation”, *The Quarterly Journal of Economics*, 111, No 2 (1996), 319-51; Bennett, Trevor, Farrington, David and Holloway, Katy, *The Effectiveness of Neighbourhood Watch*, Campbell Collaboration, no date).

The primary econometric response to simultaneity involves the use of an instrument (ie a variable which exhibits a relationship between, in this case, the number of SARs filed, but which is unrelated to rates of crime). An interesting example of the use of an instrument to overcome simultaneity emanates from Levitt (1996, above), in a study on the impact of prison population on crime. Levitt uses the status of prison overcrowding litigation to estimate prison’s impact on crime (as he finds that prison populations are influenced by the status of prison overcrowding litigation—but changes in crime rates do not appear to affect the status of litigation).

Similarly, simultaneity between SARs and crime may potentially be overcome through the use of an appropriate instrument. Some factors which likely affect the number of SARs but which are more or less insulated from changes in crime are as follows:

- the SAR reporting sector is enlarged through legislation;
- SAR reporters are finding more suspicious activity (perhaps through better systems, better feedback from LE, better training);
- reporters are filing more SARs because of defensive reporting (perhaps following a major public regulatory action);
- events encourage reporters to file more SARs (eg 9/11 or 7/7);
- reporters take the SARs regime more seriously.

However, these factors are not always easy to translate into data, and there needs to be sufficient variation in the data to tease out true relationships between variables. Whether a suitable instrument exists is not clear.

Finally, even if it is possible to determine what to measure (and to find quality data for that—not likely), and to find a suitable instrument (again, with quality data—similarly not likely), it simply may not be possible at present to estimate the impact of SARs on crime with any precision. This is largely the case because data for control variables may not exist. For example, until recently, SARs appear to have been under-utilised by most law enforcement agencies in the UK (see Fleming, Matthew H, *Asset Recovery and its Effect on Criminal Behaviour: a Review of the Literature*, London JDI, 2005 (mimeo)). As such, SARs’ crime-reducing efficacy may be spuriously related to law enforcement use or not of SARs in any analysis—unless

this is properly regulated. So for this and other activities, like smarter working and the use of advanced analytics systems in law enforcement, it will be important to control SAR data quality, SAR filing behaviour, legislation, size of the reporting sector, etc (all of which can and do change over time). At present, however, data for these control variables do not exist.

In summary, it seems fair to conclude that the time is not yet right to attempt to estimate SARs’ impact on crime. Efforts should focus on securing suitable data for one or more jurisdictions, including data for use as an instrument. In the meantime, it seems prudent, as Reuter and Truman and Reuter and Levi have suggested, to focus on intermediate indicators.

MEASURING THE SPECIFIC IMPACT OF SARs ON DETECTION

While it may not be possible at present to estimate the impact of SARs on crime, it should indeed be possible to explore the impact of SARs on ex post crime detection specifically. This brings us back—perhaps less glamorously but certainly more pragmatically—to thinking of inputs and outputs (but less of outcomes). Inputs and outputs should include the usual suspects, such as numbers of SARs received, arrests/convictions for money laundering, etc (though it is important to be mindful of the shortcomings of data on money laundering arrests/convictions noted above).

Critically, however, we will need to be smart about how we compare inputs to outputs. After all, a SAR may relate to the activity of one individual, or may refer to a series of suspicious actors simultaneously. Further, we may see multiple SARs on the same individual or individuals. With a hypothetical population of 100 SARs, there is no way to know, without getting into the data themselves, whether or not we could expect 10, 100, or even 1000 arrests/convictions from those SARs. SARs may also relate to false identities, so that even though 100 SARs may appear to refer to, say, 10 identities, it may be the case that these 10 identities are really the 10 identities of only one individual (and it is only possible to convict real people, not false identities).

The lesson here is that it is absolutely imperative to look holistically at the data. The concept should be to employ advanced analytic methods to automatically exploit all links in the data to build networks of, for example, account numbers, telephone numbers, addresses, names, and the like. This should give us an idea of what the data contain in terms of likely numbers of identities (which can be collapsed, using advanced analytic methods, to the likely numbers of individuals). This holistic, networked data view is a critical first step in understanding how inputs can relate to outputs (and is also absolutely critical in actually getting the most out of SARs in terms of actionable intelligence).

But more than simple inputs and outputs, we are probably most interested in how SARs contribute to knowledge discovery. Put simply, we would like to understand the benefit (the marginal benefit, in economics terms) of SARs in terms of the value-add of SARs in intelligence/investigations (ie do SARs provide a competitive advantage). SARs have multiple uses, which include to:

- ***Identify previously unknown criminals/terrorists, their activity, their networks.***
- Identify the previously unknown criminal/terrorist activity of known criminal/terrorist.
- Corroborate known criminal/terrorist activity.
- Enhance existing criminal/terrorist investigations (eg with new avenues).
- ***Identify/locate assets for cash seizure/forfeiture, restraint/confiscation.***
- Assist in asset recovery enforcement (and, potentially, fine enforcement).
- Prevent dissipation of assets and/or to disrupt current criminality/terrorist funding.
- Identify patterns of high- or low-volume reporting for regulatory follow-up.
- Identify the potential duplicity of regulated entities.
- ***Identify new threats/risks/trends in money laundering/terrorist financing.***
- Contribute to strategic assessments.
- Satisfy international commitments (see Fleming, 2005, above).

To measure the impact of SARs on detection specifically—to measure the marginal benefit of SARs—it seems fair to focus on the ability of SARs to tell us what we did not know before. Along these lines, we might wish to focus our pragmatic measurement efforts on three things (in bold/italics, above): SARs' ability to identify previously unknown criminals/terrorists; to identify assets for asset recovery; and to identify new threats/trends.


So given the networked SARs data (discussed above), to indicate SARs' value-add other sources of information should be added to SARs data, such as information housed in criminal intelligence databases and asset recovery databases like the UK's Joint Assets Recovery Database (JARD). The criminal intelligence and/or JARD data would be included in the networks of links between/among data. Then the real question becomes: what do SARs bring to the party? Our criminal intelligence data will relate to a number of identities/individuals (which to some extent will depend on the analytics techniques we have applied to the intelligence data); how many more identities/individuals do we know about when we add the SARs data? How many more assets do we know about when we add the SARs

data? Overall, in how many cases do SARs connect the dots and expose hitherto unknown relationships? These comparisons represent inexact science for a number of reasons—one of which is that SARs may incorrectly identify innocent activity as suspicious—but are certainly indicative of the value-add of SARs.

With similar thinking, and smart technology, we can holistically employ SARs data to tell us about money laundering and terrorist financing behaviour. In this regard, we can employ advanced unstructured text mining/natural language processing methods and tools to generate information automatically on the types of money laundering/terrorist financing which fall from the holistic view of the data. This information can then be compared, for example, to known laundering/terrorist financing typologies. As above, we can begin to think about what the data have told us that perhaps we didn't know before. We can also, again somewhat unscientifically, track the manner in which this data-driven information changes in response to changes in a country's anti-money laundering legislation (eg when new reporting sectors are added to the mix).

CONCLUSION

Measuring the impact of SARs on crime is, unfortunately, not a straightforward endeavour. Several authors have noted this to be the case, and the present paper has sought to expand on the factors which prevent us from simply "running the numbers."

For the time being at least, there is a lack of data and some complications to the underlying mechanisms through which SARs affect crime, and crime in turn affects SARs' present barriers to the application of econometric methods for estimating a policy impact. While the future may hold more promise (with better data), this suggests that a more pragmatic approach is currently required. This approach should seek to highlight the value-add which SARs bring to crime detection by leveraging a holistic, networked view of SARs' data along with other law enforcement data. 

Matthew Fleming

Senior Manager, DeticaDFI. Author's email address: matthew.fleming@deticainc.com. Opinions expressed herein are those of the author and do not necessarily represent those of the Detica Group plc or any of its subsidiaries or affiliates. This paper is based on a presentation given at the Institute of Advanced Legal Studies on March 22, 2007. The author wishes to thank John Roman of the Urban Institute and Angela Leong of the Institute of Advanced Legal Studies.