It is essential that we understand government transactions and how people interact with them if we are to improve services

Stefan Czerniawski *looks at attempts to produce figures on the amount of government transactions and highlights the need to be careful in thinking about transaction volumes.*

GDS has produced another fascinating tool, this time providing a list and volumes of government transactional services, which it turns out are used a shade under one and a half billion times a year. Richard Sargeant has a blog post introducing the endeavour, and making clear that this initial version is an alpha release.

Counting government transactions sounds as though it should be easy, but turns out to be remarkably difficult, so producing any kind of coherent picture at all is quite an achievement. There are some risks though in the apparent precision of many of the numbers which make it tempting to rely on them more than they deserve. That has the potential to be much more than a narrowly technical issue – as I was writing this post, I had coincidentally just read an impassioned piece by Declan Gaffney on the importance of getting numbers right:

> *You can't have it both ways. If the numbers are irrelevant to the issues being discussed, then they shouldn't be in play at all: if they are relevant, they merit the same level of critical scrutiny as the moral arguments being advanced. To say 'let's not get hung up on the numbers' so as to get down to the real issues is indefensible when those issues are defined in terms of what is claimed to be quantitative evidence.*

The context in which that was written is very different from the context here – but there is an underlying point which is common: if we are going to use numbers, we have a responsibility to make them the best numbers they can be.

Back then to the set of government transactions, which Richard describes as being the first time that information has been brought together. As it happens, though, it has been done before, as part of an earlier attempt to make government digital. That's not a terribly important oversight, and I mean no criticism by it, but it does mean that there are two approaches which in principle are attempting to answer the same question, and it may be instructive to compare them.

Counting transactions was the approach initially used to measure progress towards the Electronic Delivery of Government Services in 1999, with a series of reports produced which you can still find in obscure web archives. Back in 1999, there were 566 million transactions (in the Autumn 1999 report – word doc). So the first thing we can see is that transaction volumes have gone up almost three fold in the last dozen years. Which is of course nonsense, they have done no such thing, and even a fairly superficial inspection shows that different things are being counted in different ways, with inclusions and exclusions which are necessarily fairly arbitrary.

That's an important part of why the attempt to count transactions was abandoned in 2000: the approach was inherently unstable and unreliable and it was felt that it would be more meaningful and robust to count transactional services instead, which is what later reports went on to do. One reason precipitating the change, if I remember correctly, is that the advent of cattle registration meant that every cow moving from place to place would be reportable, completely swamping the overall count.

The fact that it didn't work then is, of course, no reason for assuming that it cannot work now. Apart from anything else, the motive for doing this in the first place is very different, and the core message that one and half billion is a very large number which were better made smaller does not depend on huge precision in the counting. But that still leaves us with some important issues to reflect on.

Some numbers are proper counts with real validity, some are made up (or 'estimated' to put it politely).The overall estimate can't be more reliable than its least reliable elements – so the five zeroes at the end of the 517,500,000 local government transaction are probably a fairly generous indicator of the level of precision, with the uncertainty rather swamping the ten applications for permits for burial at sea. A total presented in the way I did at the beginning of this post is completely misleading in its spurious precision (a mistake which, to be clear, Richard does not make, though it is the sum total of all the numbers in his dataset).

Even putting aside issues of arithmetical imprecision, there is a dangerous tacit assumption in this approach, that all transactions are somehow equal. They are not. Some are fairly major undertakings, some are delightfully straightforward, some may be effectively invisible to the end user because they are a by-product of something done for another purpose. So there might well be far greater value in reducing benefit claim transactions by ten minutes apiece, leaving the total number of transactions completely unchanged, than in taking the logical next step with car tax discs of abolishing them altogether. Thought about that way, a count of transactions is not, despite initial appearances, a very user focused measure. The variable we are actually interested in might be something like 'total minutes spent on government transactions which I would really rather not be doing'. Using the former as a proxy for the latter is a bit dangerous unless we are very confident that we know how they are related.

Taking that a step further, the situation is made more complicated still by the fact that transactions do not map to contacts. Some transactions involve multiple interactions. More subtly, looking up information is not a transaction, but if completing a transaction depends on that information, then from the user's perspective, it's an essential part of the journey. If we look at the transaction too narrowly, we risk falling into the trap of taking too much of a producer perspective. Even within the transaction, it may matter a lot how many stages it has, how evidence has to be provided, and how often the service design can support a single simple interaction to achieve the transactional outcome.

Sitting behind all of that, there is the question of what counts as a government transaction in the first place. GDS has attempted a definition which looks reasonable, but boils down to including any interaction which results in a change to records held by a public body, which could be interpreted very broadly indeed. If I run my company payroll and the software silently transmits relevant data to HMRC, has there been a government transaction? Is buying a bus ticket in London from a public body to be counted differently from buying one in most of the rest of the country from a commercial operator? The tool records 64.6 million library transactions a year: it's not clear what that is counting, but it doesn't appear to be either visitors (330 million) or books borrowed (300 million). To add to the slight sense of murkiness, some of the transactions counted in the tool are for the UK, some for GB and some for England.

The point of all that is not that counting transactions is a doomed endeavour or that the GDS tool is not useful. On the contrary, it has the potential to be so useful that it is important to improve on the first version (with one simple but hugely valuable improvement being a description for each line item of what has actually been counted and the provenance of the data).

The more difficult, but more important, conclusion I draw is that we need to be really careful in thinking about transaction volumes. It is essential that we understand transactions and how people interact with them if we are to improve the services of which they form a part. Transaction are not needs, so aren't the right starting point for improving service design. But making transactions smaller, simpler, and fewer is a powerful way of improving the overall service. We all have a strong interest in understanding the landscape of transactions better and the GDS tool is an important and welcome contribution to that understanding.

This article was originally published on the Public Strategist website.

Note: This article gives the views of the author, and not the position of the British Politics and Policy blog,

nor of the London School of Economics. Please read our comments policy *before posting.*

About the author

Stefan Czerniawski is a civil servant at the Department for Work and Pensions. For the last seven years, he has written Public Strategist, a personal blog on public sector delivery where this post was originally published. The views expressed are those of the author, they do not represent the position of the government.

No related posts.