

Rebuilding Macroeconomics Discussion Paper

The GFC, Systemic Legitimacy and “Rip-Off” Stories in the Daily Mail

TONY CURZON PRICE

GAVIN HASSALL

JEREMY DAVIES

JAMES COUPER

Discussion Paper

Date Published: 29. 1. 2020

<https://www.rebuildingmacroeconomics.ac.uk/publications>

National Institute of Economic and Social Research
2 Dean Trench Street, Westminster
London, SW1P 3HE

This Discussion Paper is issued under the General Call of the ESRC’s Rebuilding Macroeconomics network.

The Discussion Papers do not represent output from funded projects by Rebuilding Macroeconomics but are welcomed and relevant to our mission. These papers have not undergone a peer review process and are published to encourage discussion.

Rebuilding Macroeconomics is funded by the UK’s Economic and Social Research Council (ESRC). Our aim is to transform macroeconomics back into a policy relevant social science. We bring together scholars of from many disciplines, policy makers, representatives of civil society and interested members of the public. We are particularly interested in interdisciplinary research and bringing new methods into macroeconomic analysis.

Working Papers may include preliminary or incomplete work and they are circulated to encourage discussion and comment. Citations and use of such papers should take this position into account.

The GFC, systemic legitimacy and “rip-off” stories in the Daily Mail.

Tony Curzon Price, Gavin Hassall, Jeremy Davies and James Couper¹ 12/8/19

The legitimacy of our mixed economy capitalist model has felt less and less secure since the Global Financial Crisis (GFC). Some quite remarkable changes at the policy level seem to reflect this. For example, a UK Conservative party - the Tories - that has defined itself through free-market mantras for over 40 years reversed some of its iconic utility deregulation by deciding to introduce direct regulation of domestic energy prices. The Labour party, which had as recently as 2015 had been a proponent of such a policy (and branded “dangerous socialists” for doing so by the Tories who went on to adopt the policy), have upped the ante and called for a programme of renationalisation of some key utilities. In other words, the political system is certainly behaving as if the legitimacy of our economic system is in doubt.

But if one tries to measure legitimacy or how it has changed, it is remarkably difficult to pinpoint. The reasons for wanting to measure it are numerous: you might want to persuade nay-sayers who think that the sentiment will pass that in fact it is persistent (or vice versa); you may want to understand the processes by which legitimacy comes about in order to better design policy; you might want to better understand legitimacy in order to improve the institutions which are meant to generate legitimate outcomes; and more radically, you may want to change what is considered to be a legitimate outcome. During 2018, and after the government had legislated for that energy price cap, there was an appetite within the administration to get to grips with the legitimacy issue for these sorts of reasons.


Jurgen Habermas, in his early and influential book on legitimisation crises² is insistent that there are two dimensions to consider: whether an institutional arrangement can cope with the challenges served up by history; and whether or not it is perceived by its powerful constituent parts to be able to do so. That is: can you solve the problem, and do the right people trust your ability to do so? If neither are the case, you have a full-blown crisis. But if we are trying to operationalise these concepts, to understand how close or far we are from a crisis of legitimacy, the empirical lacunae, it seems, sits mostly in the trust and perceptions dimension. Of course, the question as to what needs to change for our system to address current challenges - the objective question - is contested and much debated. But *how* we do that, what counts as evidence, etc., is relatively settled. Not so, it seems, with the trust dimension.

Social research survey techniques make some inroads into the question. Many polling organisations run trust surveys. For example, Ipsos Mori provides these sorts of insights:

¹ This work was supported by the ESRC’s “Rebuilding Macroeconomics” initiative at NIESR, and by BEIS. Particular thanks to Angus Armstrong at NIESR for allocating the researcher time and to Luke Pereira for his interest in the analytical methods, which allowed us to use this work as a testbed for the SignalFish.io text-processing tool flow.


² Jurgen Habermas, 1973 - https://www.goodreads.com/book/show/305361.Legitimation_Crisis

Figure 1: Social Groups trusted to tell the truth, (1983-2018, Britain)



Perhaps one of the striking things in this series is how little variation there is (Figure 1). There is what looks like a long-term decline for Clergy/priests but a gently upward trend elsewhere. Bankers suffer a big dip around 2008, but their recovery has been swift and strong.

Figure 2: Measurement for trust in the British Government (1986-2019)



The results for Government are also striking (Figure 2). Trust that the government puts the nation first shows no real trend since the 2000s, and the increase in the number of people who think the system of government needs a great deal of improvement seems clearly related to a minority government blocked by Parliament in its major policy.

While there are no doubt very important truths in the relative stability of these metrics, they do not conform to the prior expectations of many politicians who believe that the climate of legitimacy has changed. Whatever the underlying truth, it would be useful to find proxies that conform more closely to the subjective experience of those whose trustworthiness is being questioned - at least it allows analysis.

Results such as these from Ipsos tend to be conversation stoppers - "why is the legitimacy of the State under such assault and what can we do about it?" is answered simply by "it isn't". Moreover, the very high level of aggregation involved makes it very hard to relate to specific policies or events in the political process - it is hard to associate changes to specific policies over which individuals or departments have agency.

We were therefore looking for a different way to measure and analyse something that could help with our understanding of legitimacy, something closer to behaviour that could be interpreted as participating in the sentiment of loss of trust. This is where the idea came to us of using measurements based on Tabloid coverage of personal finance - how are the Tabloids talking about money? How do they relate changes in the economy to what's happening in their readers' lives? And how has this changed over the years?

We provide a feel for why looking at Tabloid coverage of personal finance might offer a way into the question of legitimacy. Here is a selection of some of the more extreme front covers:

Figure 3: Newspaper front covers



This is incendiary stuff. And within the money pages of these publications, there are many more like it - here is a selection of headlines.

Table 1: Selection of Headlines

They hike loan costs, refuse insurance claims and cut savings rates... How the banks make £341 a second from YOU
Proof that energy giants are ripping us all off
Banks resort to new tricks to dodge PPI mis-selling claims as complaints top 1,000 A DAY
Britain worst for car price rip-offs
Insurers sold rip-off policies to millions - then abandoned them
Price comparison websites slammed by FCA for confusing customers
TSB customers were robbed of their life savings after IT debacle
How greedy phone giants hike our bills every 16 days
Gas suppliers using dirty tricks to hike bills
Budget airlines pocketing £265,000 a DAY in rip-off debit and credit card fees
Mortgage rate rip-off: banks stand accused

And here is a word-cloud made of all rip-off article headlines (Figure 4). "Fraud", "complaints", "rip-off", "mis-selling" and "scam" are all prominent - and that is just the headlines. We used machine learning and text analytics to probe trends and patterns in this sort of coverage more deeply. We built a database of every money-section article published (and reachable online) in the online Daily Mail since 2000 - that was 240,000 articles. We trained a machine learning algorithm to spot articles about rip-offs. We used sentiment analysis and other text metrics to determine the emotional stridency and valence of these articles. And we used automatic classification tools to help us map the article subjects.

[illegible]


Descriptive statistics

Figure 5:



There was a fall in the total number of money-section articles published online between 2000 and 2006, and an increase between 2008-2012 (Figure 5b). The absolute number of rip-off articles has been much less variable than the total number of money articles (Figure 6a). Therefore, the relative preponderance of rip-off articles has increased since 2012.


Figure 6:



An overlay of the rip-off article frequency and real per capita GDP is interesting in showing the absence of any obvious relationship - except that perhaps the immediate post GFC spike in articles corresponds to the period of fall in real GDP per capita (Figure 6b). The early period is characterised by rising GDP per capita and rising rip-off articles, and the later period by rising GDP and high but stable levels of rip-off articles.

The sentiments associated with rip-off stories are always negative. We use the Google emotion analysis engine which returns a value between -1 and 1 to describe the degree of negativity or positivity of the emotional load of an article (Figure 7a). All our rip-off stories have negative emotional score and are systematically below the non-rip off articles. It should be noted that Google reports that long complex pieces of text, as most of these articles are, can give imprecise measures of emotional valence if the text contains mixed messages. For example, an outrage followed by a suggestion for its resolution, can show a higher emotional valence than a mild outrage alone.


Figure 7:



The stridency of the stories - how *loudly* they convey their message, the volume on the sentiment dial (which Google measures as a positive scalar) - is relatively constant until 2012, and then rises steadily (Figure 7b). This rise is more pronounced for rip-off stories than for non-rip-off stories.

Using text classification tools to automatically tag articles, we can estimate the dominant subject of each story. We find that Finance regularly tops the list in terms of rip-off frequency - both before and after the GFC. However, as the stream version of this diagram (Figure 8) makes clear, by 2010, Finance and Banking were coming off their rip-off peaks. The fall in total rip-off frequency stalls by 2014 and starts to rise again in 2017. The increase comes from “energy”, “autos”, a broad, composite “other” and “telecoms”. So despite the remarkable dominance of banking, loans and finance, we find that rip-off stories move from subject to subject. “Real estate” and “mortgages”, the subjects perhaps most closely linked to the macro-causes of the GFC, show a low and relatively constant level of “rip-off” stories.


Figure 8:



We see that over this period, “rip-off” stories are increasingly likely to contain “household name” firms. From 2008, the total number of well-known firms in the “rip-off” corpus is markedly higher than before, and the maximum level is reached post 2017 (Figure 9a).

A slightly different measure - the average number of well-known firms mentioned per rip-off article, rises from 2004 and increases strongly after 2017 (Figure 9b).

Figure 9:



We postulate that the person whose beliefs are substantially influenced by coverage in the press will increasingly view the economic system as being rotten - as lacking in legitimacy - as the following increase:


- The volume of rip-off stories
- The number of sectors they touch
- The frequency with which household names are implicated
- The stridency of the stories
- Their degree of negativity (although this seems sensitive to the problem of mixed messages within a single article)³



If we look at all five measures - for example in the accompanying gap minder animation⁴ - we see that the levels of these measures are low but growing in the early 2000's, that they rise markedly during the crisis, especially 2009-2010, and that levels of de-legitimation are at a persistently high level thereafter, with new subjects beyond finance growing strongly in the later period. The size of the dots represents article frequency, and their color represents degree of negativity. The y-axis represents stridency, and the x-axis represents the mention of well-known firms. Our hypothesis is that the redder, the bigger and the further to the top right are spots, the greater the sense of de-legitimation.

³ Google, in their notes on their sentiment scoring algorithm, note that long texts with strongly negative and strongly positive emotions - for example a combination of problem and its solution - are hard to score on "degree of negativity", as the measure tends to average out the extremes. It is interesting that a human reading of rip-off stories appears to suggest that solutions, like switching services, are often part of rip-off stories in the later period. We conjecture that this may be related to newspaper business models, since the referral fees on switches can be very high - of the order of £50 per completion.

⁴ This can be found online at <https://www.youtube.com/watch?v=xfF7oOSovFI>

Figure 10: Development of 'De-Legitimation': 2003 - 2018





Analysis and hypotheses

There are a number of possible hypotheses one could advance about the peak in rip-off stories between 2008-12, and the subsequent levelling, for example:

1. The number of rip-offs rose (perhaps because firms were looking for new sources of profit)
2. The tolerance towards rip-offs fell as belts tightened - which we might call the Galbraith *bezzle* hypothesis⁵ - leading to interest in their exposure
3. The profitability of rip-off stories to publications rose because of changes in their business environment
4. The demand of readers for scapegoats onto whom to focus anger increased, and the Mail was responsive to reader demand


Is there "narrative independence"?

The first hypothesis is the one that questions narrative independence. In order to actually answer the question, we need some independent measures of the frequency of the rip-offs that the stories cover. This is often painstaking data to collect, although regulatory probes can be a source of very rich data. For example, the CMA energy market inquiry measured in great detail the consumer detriment from overcharging in that sector - the amount paid over a competitive benchmark by inert, loyal customers. The two graphs below show the CMA's measure of the rip-off and the frequency of energy rip-off stories. Rip-offs are present and large in all the years the CMA measured it, while the story frequency and intensity varies substantially. This is suggestive of narrative independence. Conversation with executives and industry-insiders tends to reinforce that conclusion - in both finance and energy, informal conversations with executives suggest a level of surprise at the attention now being paid to business practices such as "sweating the back-book" (a finance sector expression for what has come to be called "the loyalty penalty" in regulatory circles). The sense is that this is what these businesses have always done, and that complaints about their business models are what is new.

⁵ In his *Great Crash of 1929*, Galbraith describes the following process: "At any given time there exists an inventory of undiscovered embezzlement in — or more precisely not in — the country's businesses and banks. This inventory — it should perhaps be called the *bezzle* — amounts at any moment to many millions of dollars. It also varies in size with the business cycle. In good times people are relaxed, trusting, and money is plentiful. But even though money is plentiful, there are always many people who need more. Under these circumstances the rate of embezzlement grows, the rate of discovery falls off, and the *bezzle* increases rapidly. In depression all this is reversed. Money is watched with a narrow, suspicious eye. The man who handles it is assumed to be dishonest until he proves himself otherwise. Audits are penetrating and meticulous. Commercial morality is enormously improved. The *bezzle* shrinks."

Figure 11:

Average SVT price (based on the annual bill for a dual fuel direct debit typical consumption) and a forward-looking industry-level benchmark of direct costs




Source: CMA analysis of data collected from the Six Large Energy Firms, Ofgem and ICIS.

The vertical distance between the blue and the green curves is a good approximation to the size of the energy overcharging. In this CMA graph is a first approximation to the presence of "rip-offs" it shows the degree to which people are paying very different prices for essentially identical products.

The graph below shows the frequency of energy rip-off stories. It has a 2008 peak, rises again in 2011 and then slowly falls. Comparing it to the objective underlying degree of overcharging in the graph above, there is no obvious relation. If anything, the volume of stories is falling after 2012 when the overcharging harm is, if anything, growing.

Figure 12:



Although our data is suggestive rather than conclusive, we tentatively tend towards the view that there is a substantial degree of narrative independence. In any case, we want to highlight the degree to which these new measurement techniques hold out the promise of exploring these questions empirically in a way that has been hard to do before.⁶

The Galbraith Bezzle Hypothesis

J K Galbraith's bezzle hypothesis seems like a natural one and suggests an explanation for some of the patterns found that is not very far from economic orthodoxy.⁷ The GFC lowers incomes, and as belts tighten, people look around (assisted by the Daily Mail) for ways to alleviate the pain. These stories are responding to objective economic demand for information. This pattern is all a part of markets for actionable information, and the frequency pattern could be seen as an indicator of these markets for information working well.

However, this explanation does not fit all the salient facts about rip-offs. First, we know from analysing these markets in detail that mostly, people do not take action in the light of these rip-off stories – the Competition and Markets Authority's (CMA) energy market investigation showed that two thirds of households, often the poorer ones, did not exercise the agency they could in the face of overcharging. Follow-up work by Citizens Advice and the CMA on the "loyalty penalty"⁸

⁶ We would like to note that these questions and hypotheses are the bread and butter of media sociology and anthropology. We think that this sort of empirical analysis has a great deal to learn from the groundwork and theory developed by these disciplines.

⁷ Galbraith is characteristically subtle in his description of the effect - he hints at the moral outrage involved, for example. We use this label as a respectful shorthand, not as a complete description.

⁸ <https://www.citizensadvice.org.uk/about-us/our-campaigns/all-our-current-campaigns/citizens-advice-super-complaint-on-the-loyalty-penalty/>

suggests such “rip-offs” are common to many sectors, and consumer inaction in the face of them is similarly common, so it seems unlikely that the stories are produced as part of well-functioning information markets.

Another reason for doubting this reductive version of Galbraith's bezzle is that the stridency index and the frequency of household name mentions rises well after the height of the crisis. In other words, a simple mechanism relating to objective economic facts post-GFC does not seem to explain the pattern we find.

The "clickbait" hypothesis

Here is an alternative hypothesis. The GFC creates anger, and the Mail's business model is to attach stories to that anger: readers expend more attention on stories that provide an outlet for that anger. This attention is monetisable through advertising, perhaps even advertising of solutions to the problems being written about, like the very rich pickings from referral fees on switching services. Editors of money pages can increase the number of such articles without seeing a diminution of total clicks, perhaps because the anger does not abate - there is not a fixed amount of anger to be allocated across stories; rather, this kind of story can compete for attention against other types of story. Moreover, journalists find stories of rip-offs (that had previously not been uncovered, or that had but had failed to gain clicks, perhaps because of the Galbraith bezzle hypothesis) in response to their editors' demands for more of these stories in the competition for attention.

A person whose world is constructed as a recipient of this process of attention competition will consider more and more of the economy to be part of “rip-off Britain”. It is easy to imagine that this has an accelerator effect - another story about rip-offs increases a sense of indignation, because it becomes another sign of a *generalised* malaise. A genre once concentrated mostly in finance products spreads to other sectors. When the democratic process fails to respond to the objective (but longstanding) problems that are uncovered, perhaps the stridency of the reporting rises.


But this account of the late rise of stridency is probably not the only - or perhaps not the most plausible - hypothesis. During this period, we not only see austerity and the GFC, but also the complete transformation of newspaper business models. The iPhone enters our lives in the same year that Northern Rock leaves it. Competition for attention leads to increasing quantities of “click bait”.⁹ This adds to the increased demand for rip-off stories.

⁹ I am very thankful to Robert Cottrell, editor of TheBrowser.com and onetime journalist himself, for the following account of the rise of clickbait. Why, I asked him, did editors not make headlines irresistibly enticing before the ad-driven online model? Robert's answer is that in the print world, once a reader has bought your paper, the reader and the editor are in a cooperative game - their interests are aligned: the reader does not want to waste attention; the editor wants the reader to repeat their purchase tomorrow. However, online, the editor does not make advertising money until the reader is on the page. Therefore, the headline is a “come-on”, a sales pitch, and not a move in a cooperative information-processing game. In the print world, the front page may well share the qualities of click-bait - indeed, see the resources and effort lavished on the front page of a print media publication - but the inner pages are not.

This overall account - the monetisation of anger - would be consistent with our descriptive statistics. But it hypothesises a quite complex set of micro-causal mechanisms for which we do not (yet) have strong evidence. However, we believe that the data set that we have assembled, combined with the enrichment from machine learning and text-metrics, offers the prospect of rich empirical examination of this sort of account.

For example, our hypothesis would imply a correlation between stridency and the shift of advertising revenue to online media. The *bezzle hypothesis* would not predict any such correlation. Indicatively, but, of course, far from conclusively, the graph below shows a close correlation.

Figure 13:



Our hypothesis posits a causal relationship based in newsroom practice between the types of stories that editors express a desire for and click-through rates. This sort of micro-causal mechanism can be lent plausibility by direct empirical observation of changes in newsrooms as online practices were integrated into them.¹⁰ Data about actual article click-throughs and

¹⁰ For example Rojas (2018) provides a detailed anthropological study of the transformation of the newsdesk in a major Chilean newspaper as online teams were integrated.

profitability - data which exists inside the news organisation - could in principle be used to test the micro-causal account further.

Research methods and questions similar to the one described have been relatively common in media sociology and media studies, and we believe that there is much that policy professionals (and economists especially) could learn from that literature. Perhaps the allure of large and rich data sets, with the opportunity for a great deal of bread-and-butter econometrics, will tempt colleagues to look again at important insights from these neighbouring disciplines.

The Scapegoating hypothesis

The Scapegoating hypothesis, our fourth, is an attempt at explaining why *these* are the sorts of stories that have enjoyed an increase in supply and demand. After all, the clickbait hypothesis might have predicted an increase in the prevalence of cat stories rather than rip-off stories. The scapegoat hypothesis, like the simple version of the bezzle hypothesis, tries to explain the *content* of the stories: why is there a demand for this sort of product?

The scapegoat hypothesis was suggested by the historian Simon Schama in a 2010 [essay for the FT](#)¹¹ that seems remarkably prescient from a policy perspective.

He writes:

Historians will tell you there is often a time-lag between the onset of economic disaster and the accumulation of social fury [Act one is...] Act two is trickier. Objectively, economic conditions might be improving, but perceptions are everything and a breathing space gives room for a dangerously alienated public to take stock of the brutal interruption of their rising expectations. What happened to the march of income, the acquisition of property, the truism that the next generation will live better than the last? The full impact of the overthrow of these assumptions sinks in and engenders a sense of grievance that “Someone Else” must have engineered the common misfortune. The stock epithet the French Revolution gave to the financiers who were blamed for disaster was “rich egoists”. Our own plutocrats may not be headed for the tumbrils but the fact that financial catastrophe, with its effect on the “real” economy, came about through obscure transactions designed to do nothing except produce short-term profit aggravates a sense of social betrayal. At this point, damage-control means pillorying the perpetrators: bringing them to book and extracting statements of contrition. This is why the psychological impact of financial regulation is almost as critical as its institutional prophylactics.

<https://www.universiteitleiden.nl/en/research/research-projects/social-and-behavioural-sciences/the-infrastructure-of-news-newsroom-ethnography-in-chile>

¹¹ <https://www.ft.com/content/4526d52c-6506-11df-b648-00144feab49a>

Schama sees in 2010 that it remains essential to deal with the unspent social fury of 2008, caused by the expectations betrayed by the GFC. In the essay, he emphasises the importance for legitimacy that justice be seen to be done, and the importance of objects of hatred on whom to pin the blame. The continuing appetite that we have measured for rip-off stories, for stories that implicate more and more of the “system”, seems consistent not only with the view that justice was *not* seen to be done, but also with the search for objects onto whom to attach the social fury.

It is tempting to see the rip-off coverage in papers like the Mail in terms of a scapegoating logic: their intention is not to start with anger and solve a problem, as the bezzle hypothesis might have it; but rather to identify possible scapegoats and go through the motions of justice, if not its actual execution. This might explain, for example, the fact that our delegitimising markers seem to bring in more and more subjects outside finance. It might also explain the lack of remedial action being taken by readers: they are perhaps not asking for a solution to the apparent problem posed by the article, and the placing of the feeling of indignation onto an object may in fact be what they are demanding. The increasing and persistent dispersion in the bubble animations noted above are consistent with the scapegoating hypothesis - new objects for the attachment of indignation are sought. It might also explain the fact that the stories do not appear to be objectively linked to the sectors and markets that caused the GFC - mortgages and real-estate have a persistently low and small score on our various delegitimising metrics.

Perhaps we can go further still and say something about the role of an institution like the Daily Mail in this social process of scapegoating: is it part of the healing, or part of the wounding?

Even once the immediate financial meltdown threatened by the GFC was averted by policy response, there was still plenty of room for a sense of betrayal, even for asset-owning middle classes. Their long term future prospects generally looked much less rosy post 2008 than pre: the *great moderation* was good for them. But why direct anger towards retail products and sectors that were not particularly implicated in the GFC?

It is worth remembering the origins of the scapegoat, from Leviticus 16. There was bad blood amongst the Israelites because two sons of Aaron had mysteriously died in the Temple. Aaron had good reason to be angry. He was also, as the brother-in-law of Moses, extremely powerful. The Lord spoke through Moses and ordered that Aaron should contribute a ram to be sacrificed at the temple. Sacrifices of this sort were usually shared amongst the community as cooked meat. But seeing that this would not end the anger, The Lord ordered also that a goat be found and taken to the desert. No one would benefit from its death. There would be no meat from temple sacrifices. It would be a permanent subtraction from the collective wealth. But it would carry with it the anger - it would be a valued asset killed and wasted in order to drain that anger.

Is the Daily Mail in its rip-off stories performing a scapegoating function? (This was certainly the view of one energy company CEO when one of us suggested to him that once they had lost the support of the Tabloids, the price cap was almost inevitable). We do want to emphasise that even if the Mail performs a scapegoating function, the tabloids brought to the agenda a whole lot of actual bad business practices and have added hugely to the pressure that they be addressed

through the democratic process. In other words, the Daily Mail process is not as wasteful as the biblical scapegoat. The scapegoat of Leviticus subtracts from the material wealth of the community; but the scapegoats of the Daily Mail, to the extent that they lead to policy response or behaviour change, contribute to allocative efficiency and so add to the common wealth.

Nevertheless, there is one way in which the Mail might be thought to be similar to the scapegoat: lancing the boil of social anger by attaching it to a sacrificial goat might be seen as a way of obscuring the actual causes of the problem. Why did Aaron's sons die when they got too close to the Holy of Holies? Was it in order to preserve the mystery and authority of the Ark of the Covenant that they had to die? So wasn't the goat a victim of the fight to protect prevailing power structures? Similarly, the attachment of anger to "fat cats" in regulated sectors - however reprehensible their behaviour might be - could ultimately be the decoy helping to save the wider interests of asset-owners.

Use in policy evaluation

What is the policy-analyst and the advisor to do with these tools and these sorts of results? I want to mention two types of uses: first, the rather classic purpose of tracking and evaluating policy success; and second, the perhaps more interesting, and certainly more radical, idea of evaluating and influencing the institutions and mechanisms by which legitimacy is produced.

Let me provide a very concrete example of the first.

Andrew Lord Tyrie, the recently appointed head of the Competition and Markets Authority (the UK's antitrust authority) told the journalist Will Hutton in a recent podcast¹² that he saw the role of his institution as addressing the crisis of legitimacy in capitalism. In the interview, Tyrie argues that the harm caused by anti-competitive behaviour goes well beyond the immediate harm that is analysed and remedied through antitrust tools: "I have been watching [...] an erosion of confidence in free enterprise, in capitalism [...] In this job now, one of the tasks is to ensure that erosion of trust does not spill over into a crisis of legitimacy ... the discontent is real ... people feel they're being ripped-off ... I don't believe doing nothing is a sustainable option". In other words, the problem is not just, or perhaps mainly, the immediate harm caused by the rip-offs: it is the harm the rip-offs cause to "free enterprise and capitalism".


If that is indeed to be the role of one of our most powerful economic institutions, then we need to be able to measure and track changes in legitimacy. The methodology presented here is clearly a starting point. Moreover, if it should be confirmed that the scapegoating logic is such that anger hops from one sector to another, the project of "saving capitalism" by incrementally solving "rip-off" markets might be questioned - and that is however much one approves of incrementally solving rip-off markets.

We noted the importance of the energy retail rip-off and, eventually, the price cap as a policy response. We can track the relationship between sentiment, policy and politics using our

¹² "Goliath" - <https://www.bbc.co.uk/programmes/m0003zvn> , 14/4/19, BBC Radio 4

database. For example, in the graph below, we show the major milestones of policy against frequency and sentiment of energy rip-off stories. We see a build-up in the frequency and negativity of energy rip-off stories up until Ofgem's retail market review final report in early 2013. This seems to abate after the Milliband and Cameron Parliamentary interventions of 2013. There is a flurry of political activity all the way through to the announcement of the price cap in 2017, and this corresponds to a period of relative moderation in frequency and negativity. This falls to very low levels for three quarters up until the implementation of the price cap in January 2018. After that, however, frequency levels are back to historically high levels, even if negativity remains subdued. It is certainly too early to tell whether the price cap has had more than a fleeting effect on the atmosphere of legitimacy, whatever good it has objectively done, especially to poorer customers.


Figure 14:



A smoothed version of the same data (Figure 12) makes the point even more starkly - it looks as if article frequency falls as long as political attention is given to the issue, but once policy “resolution” is achieved, it is possible that the frequency of de-legitimation simply returns.

Finally, if we look at the progress of the energy rip-off story in the context of all other rip-off stories in the graph below, there is *some* indication that improvements in energy are offset by a deterioration in other subjects - telecoms and autos in particular. In other words, if the hope, as expressed by Lord Tyrie, is to address a systemic legitimacy issue through solutions to actual and specific problems, this sort of tracking should be useful in determining whether this is working.

Frequency of Rip-off Articles by Subject



Using these methods for institutional design

We now discuss the second way in which this sort of analysis might be helpful to policy: the attempt to improve the institutional mechanisms for the production of legitimacy.

The hypotheses posit a very extreme and hypothetical agent, one whose beliefs can be proxied by a database of DailyMail articles. (The reason it was the Daily Mail happened mainly for pragmatic reasons - that the Mail has the best online archive of all the Tabloids. But nevertheless, the Mail has the largest circulation online of the Tabloids, and it was credited with being particularly influential under the May premiership)¹³.

The way in which our general beliefs and attitudes are formed - including our attitudes to our collective economic arrangements - is a highly contested question in economics. It is even contested whether economics needs to know anything about belief-formation beyond the axioms of rational choice theory. One very stark and, in regulatory policy circles, influential version of the required assumptions about belief-and-attitude-formation comes from the Chicago school, for example Stigler and Becker's 1977 formulation in "De gustibus non est disputandum". The explicit goal of that article is to demonstrate the richness of the assumption that high-level household

¹³ See, for example <https://www.theguardian.com/media/2017/mar/30/paul-dacre-theresa-may-private-dinner-daily-mail-editor-no-10> as well as the more polemical <https://www.opendemocracy.net/en/opendemocracyuk/may-was-daily-mails-pm-day-after-it-told-her-go-she-went/>

preferences are fixed, and that a great deal of behaviour can be explained by rational choice combined with changes in prices and incomes. The assumption of this sort of well-behaved "household production function" has been hugely influential in policy circles and lies at the heart of a great deal of the normative power of micro-economics.

The assumption implicit in hypotheses three and four, and to some extent in the bezzle hypothesis, is that world-views, beliefs, attitudes and emotions are not fixed. Instead, we consider a hypothetical agent represented by the beliefs in one particular publication. Stigler and Becker's 1977 argument is epistemologically very interesting - they do not pretend to demonstrate that preferences are stable and beliefs are rationally derived; instead, they want to persuade economists that this is a rich assumption that generates insights. They offer that the alternative is a sort of methodological dead-end. If tastes are variable, then who is to do the scientific work of accounting for them - "psychologists? anthropologists? phrenologists? Sociobiologists?", they ask. The disdain is clear: since these disciplines clearly should not be involved, economics had better show that it has a rich and methodologically creative model of behaviour. They propose an extreme version of homo economicus to fulfil this role.

Our implicit model of attitude, taste and belief production is as unrealistic as Stigler and Becker's. Some aspects of it are subject to empirical testing. For example, Foos and Bischof (2019) have produced evidence that Liverpool's 30-year boycott of The Sun accounts for differences in Brexit attitudes and voting patterns in that city.¹⁴ But most importantly, just as Stigler and Becker wanted to demonstrate that their assumption was fruitful, we believe that the development of these new tools to measure narratives mean that alternative assumptions about belief formation are now fruitful for policy and for research.

We believe that our beliefs, the narratives we are exposed to, the conversations we have about them, the way in which we process them and make them our own, are immensely important in shaping social outcomes. They go through many channels and mechanisms, most of them uncharted. But the work presented here shows that quantitative analysis of those stories and the mechanisms of belief formation is far from hopeless.

We hope to have shown that an institution like the Daily Mail is an important element of that belief-formation process, and we have suggested micro-causal accounts of content-formation whose influence on beliefs about legitimacy can be analysed. For example, we have suggested that the rise in stridency of rip-off stories may have been influenced by the change in media business models and subsequent changes in editor/journalist incentives.

These mechanisms of belief formation, of course, can be influenced by various policy levers. For example, Ofcom already has a duty to ensure the country enjoys media plurality, and has some powers to attain this goal. Ofcom's framing of that goal¹⁵ includes "Preventing any one media

¹⁴ http://www.florianfoos.net/resources/Foos_Bischof_Hillsborough_APSA.pdf

¹⁵

https://www.ofcom.org.uk/data/assets/pdf_file/0024/84174/measurement_framework_for_media_plurality_statement.pdf

owner, or voice, having too much influence over public opinion and the political agenda.” In terms of the tools and powers it has, it says:

Ofcom is required, in carrying out its principal duty, to secure various ends, including the maintenance of a sufficient plurality of providers of TV and radio services. Regarding media plurality, Ofcom also has specific statutory duties which relate to the periodic review of the media ownership rules and to report on whether certain media mergers operate against the public interest (at the Secretary of State’s discretion).

In other words, public policy is already involved in assessing these mechanisms of belief-formation. The analytical methods presented here should enrich the measurement of “voice” and could suggest new policy interventions. For example, the definition of a “voice” that could have undue prominence in the formation of public opinion, might include the voices of groups, or perspectives, rather than just the voices expressed through specific channels. The methodologies used in this paper offer natural ways to quantify “voice” and even ways to measure changes to that other elusive quantity that Ofgem is required to be guided by - “influence over public opinion.”

Appendix A: Methodology

1. Machine learning methods

To identify rip-off articles in the Daily Mail's money section, the machine learning and web scraping tools of SignalFish were used. SignalFish collected the entire corpus of Daily Mail articles that were live on their website, from which a subset was gathered for developing a machine learning model. This required us to identify "positive" and "negative" examples of articles that would train a model to identify similar articles to the set of positives and jointly dissimilar to the set of negatives. We also had to identify what pieces of information from the articles we wanted to use for training the model, i.e. the headline, summary, and the full text from the article.

The rough criteria used for defining a rip-off article is summarized below:


A rip-off article is one that brings the legitimacy of firms into question. By legitimacy, we mean business practices that generally treat customers unfairly, e.g. hiking prices, using fine-print in contracts to benefit firms over customers, or simply not providing the service paid for. An important consideration in classifying an article is in how the article frames the topic. For example, just reporting on interest rates rising would not count, but emphasizing how banks rip-off consumers by raising borrowing rates whilst keeping saving rates the same following an increase in the bank rate, would count.

We began with a training set of 100 positive (rip-off) articles post-2018 from the Daily Mail and 150 negative articles that were deemed definitely not rip-off articles. The model was first trained using just the headlines and summaries for scoring the headlines and summaries of other articles based on how similar they are to the positives and dissimilar to the negatives. A score defined over $[0,1]$ was given to each article, where the closer the score was to 1 the more similar the article was to the rip-off articles. The model was used to classify 500 random Daily Mail articles published in 2019 and we found it had some ability to consistently identify rip-off articles. This is because when the articles were listed from highest score to lowest score, the first 23 articles were all rip-off articles – coinciding with articles scored above 0.85. However, there were still rip-off articles appearing lower down the scoring and mixed in with non-rip-off articles based off their scores. To push types of rip-off articles that may not have been strongly present in the first training data higher up the scoring and the non-rip-off articles lower down the scoring, we added to the training data from this exercise. The obvious rip-off articles were added to the set of positive training data, and the negative articles whose score was mixed in with the positive articles were added to the set of negative training data. This technique proved to be an efficient way of identifying more positive and negative articles for increasing the size of the training data – which for improving model accuracy, a larger training set is preferable.

On the next model iteration, we compared how two models scored the same set of articles: one trained with headline and summary data, and the other with headline, summary and the full content of the article. We found that for identifying rip-off articles, using the full content improved our ability to identify this type of article. This is because we were finding a clear clustering of actual rip-off articles in the top of the scoring, versus the headline and summary case. We once again

added the clear highly scored rip-off articles, and clear highly scored negative articles to the training data.

An issue that arose was the model was having some difficulty in accurately scoring the rip-off articles pre-2015. This was likely due to the initial training data being primarily from post-2018, and as the language of rip-off articles changed, the signal we were identifying from earlier articles was getting weaker. To overcome this, we added more rip-off articles over time to the training data. This process was considerably sped up by having our most recent model at that time score 500 articles per year over 1999-2019, and further positive and negative articles were selected from these.¹⁶ We ended up with a final training dataset of 637 positive articles and 930 negative articles. The model that was subsequently trained was used to classify another random sample of 500 articles per year from 1999 to 2019. It was found that expanding the training data over the full time dimension of our sample considerably improved the model's ability to identify rip-off articles. This is because when the articles were listed by their score, the highest scored articles were almost all rip-off articles. Using our final "MLTny" model to classify the entire Daily Mail's money section resulted Figure A1 showing the frequency of the article scores. An indication that the model is able to discern noticeable differences in articles is when there is a clustering of article frequency at either ends of the [0,1] interval, which is shown in Figure A1.



To impose a binary classification of whether an article is or is not a rip-off article, we had to decide upon a cut-off score. This was done by starting at a particular score, e.g. 0.85, and looking at the fifty articles scored immediately below this threshold. If over 25 of the 50 articles were not rip-off articles, then the threshold was moved up to 0.86 and repeated. The score that we ended up with

¹⁶ With thanks to Emilia Perry-Poletti and Katie Schöffmann for helping classify the many articles we had to deal with and expanding our training data. They also helped confirm the validity of how the initial researchers were subjectively defining rip-off articles. This is because their independent assessments of what were or not were rip-off articles, from being provided the rip-off article definition above, strongly correlated with each other and others.

was 0.93 where 28% of the immediately 50 lower-scored articles would not be considered a rip-off article. In contrast, by taking a sample of 100 random articles from the final classification used, 87 were rip-off articles – i.e. an inaccuracy rate of 13%. The likelihood that a randomly selected article is a true rip-off article does fall as we move further down the scoring, hence why the worst quality is seen at the bottom of the threshold.

The reason why a binary score was imposed was that the gradient of the scores above 0.93 did not seemingly translate into a measure of the extent to which articles were rip-off articles, in terms of, say, the strength of condemnation of the actions of firms. Likewise, two articles with a score of 0.5 which may be decisively not rip-off articles would not indicate the same presence of a rip-off narrative than one genuine rip-off article with a score on 1 would have. For this reason, the machine learning tools provided an efficient way for classifying articles into either rip-off or non-rip-off articles based off the threshold score identified.

2. Further data analysis, cleaning, and Neo4j

We were able to handle large volumes of data in such an efficient manner due to the support of SignalFish who created us a Neo4j graph database for data querying. Above the ML scoring, they were able to extract from all the articles: entities, classifications, article-level and sentence-level sentiment analysis. Entities were obtained from inspecting the given text for known entities (proper nouns such as public figures, landmarks, etc.), and returning information about those entities – and context-specific sentiment analysis was applied to each entity identified. Articles could contain just a few entities to hundreds of entities depending on their length and content. The sentiment analysis had two components: a sentiment score, and magnitude score. The sentiment score referring to whether the text or entity displays a positive or negative emotional tone, and the magnitude tells you how much emotional content is present in the text or surrounding the entity. Both scores were obtained for all entities, the overall article, and the sentence-level. The Entity type gave six broad categories that an entity could fall into – the most useful for our purposes was “Organization” entity type. The classification was drawn from a predefined set of 640 categories that the article text most closely resembles.

The metagraph for the Neo4j database is shown in Figure 2, indicating the general structure of the database and how each article was dissected.




Figure 2: Neo4j Meta-graph.

The main benefit of working with Neo4j in this use of big data was its speed of performing search queries. Our relatively small database was 5 Gigabytes containing 244,211 articles and 2,199,935 unique nodes (the circles in figure two) with 12,527,736 total nodes used over all articles (as some nodes are related to more than one article). For example, the “Article” node contains information on the publication date, article headline, summary and content, overall document sentiment scores, the number of categories it relates to, and any other information we may wish to assign to this node. Standard search queries and manipulations of the database are typically completed in under a second, and even the most complicated queries with hundreds of constraints typically return in under a minute.

2.1 Classification categories

One of the manipulations of the initial graph created by Signalfish was conducted to measure shifts in the focus of rip-off articles. Each article was identified as relating to some classifications, and provided a score in [0,1] on how relevant each classification was to the article. There were 640 possible classifications, for example, “/Finance/Credit & Lending/Credit Cards” which would suggest a body of text relates specifically to credit cards, which is a part of credit and lending, which is part of finance. Some articles were assigned to several categories, but most to just one; the frequency of categories each rip-off article was assigned to is shown in Table A1.

Number of Classifications (x)	Number of Rip-off Articles with x No. of Classifications
1	11976

2	4878
3	940
4	73
5	12
6	4

Table A1. Frequencies of rip-off articles with either 1, 2, ..., or 6 categories assigned to them.

For articles with more than one category, the category with the highest “weight” was used in the final analysis. This is because some articles appeared in, say, insurance-related and automobile-related articles if they were related to car insurance. Whichever of these themes had the dominant narrative in the article – reflected by the weight – would be used to classify the article as its dominant category.

We were left with 294 dominant categories in rip-off articles, 263 of which had fewer than 100 articles in each over 1998 to 2019, representing 16.4% of all rip-off articles. This indicates 83.6% of rip-off articles were contained in 31 categories. To analyse the focus of rip-off articles over time, groups of categories that related to similar topics were aggregated into 15 categories: autos, credit cards, transport, mortgages, banking, energy, shopping, telecoms, travel, loans, real estate, finance, insurance, pensions, and “other.”

For example, the autos category was comprised of all articles from the categories:

"/Autos & Vehicles/Vehicle Codes & Driving Laws", "/Autos & Vehicles/Vehicle Parts & Services/Vehicle Parts & Accessories", "/Autos & Vehicles/Motor Vehicles (By Type)/Motorcycles", "/Autos & Vehicles", "/Autos & Vehicles/Motor Vehicles (By Type)/Trucks & SUVs", "/Autos & Vehicles/Vehicle Parts & Services/Vehicle Repair & Maintenance", "/Autos & Vehicles/Vehicle Shopping/Used Vehicles", "/Autos & Vehicles/Vehicle Parts & Services", "/Autos & Vehicles/Motor Vehicles (By Type)", "/Autos & Vehicles/Vehicle Parts & Services/Gas Prices & Vehicle Fueling", "/Autos & Vehicles/Vehicle Shopping", "/Autos & Vehicles/Motor Vehicles (By Type)/Hybrid & Alternative Vehicles", "/Autos & Vehicles/Classic Vehicles", "/Autos &

Vehicles/Vehicle Codes & Driving Laws/Vehicle Licensing & Registration", and "/Business & Industrial/Automotive Industry".

The “other” category was comprised of 138 categories, generally related to: arts and entertainment, television, night-life, healthcare, pharmaceuticals, hospitals, dentistry, vitamins and supplements, nursing and caring, home improvements, home and garden, beauty, fitness, fashion and style, restaurants, education, electronics, computing, gaming, gambling, performing arts, legal services, music, sports, sports teams/clubs, pets and animals, veterinary, movies, e-commerce, and construction. These categories on their own were not particularly dominant over time, some with only tens of mentions over twenty years, and is why they were grouped together. The growth of this category in rip-off articles post-2015 may indicate a widening scope of industries being targeted.

2.2 Well-known firms

The extent to which rip-off articles mention “well-known firms” may be an important indication of their effect in delegitimizing the market in the eyes of readers. In contrary to small-scale scams from tradespeople or second-hand car dealers, when well-known firms are portrayed as the ones responsible for the rip-offs, this may influence people’s attitudes towards the market more generally.

The measure of well-known firms is based off the entity extraction, and most firms mentioned in articles – particularly the well-known ones – will be identified through an entity node for the article. The “Entity type” will also identify them as an “Organization”. The list of well-known firms was created by seeing which organization entities were the most frequently mentioned in the Daily Mail’s money section and selecting the ones that counted as “well-known.” Firms with fewer than 15 mentions over 1998-2019 were not included. The full list of 195 firms used is shown in Table A2. This list was also cross-referenced with a list of the 150 biggest UK brands, and almost all of those firms were contained in this list.

aa insurance	cadbury’s	halfords	nestle	southern energy
abbey national	canada life	halifax	netflix	sse
admiral	capital one	hargreaves lansdown	network rail	standard chartered
aegon	champions league	hbos	new look	standard life
air france	channel 4	hmv	nissan	starbucks
aldi	cheltenham & gloucester	homebase	northern rock	superdrug

alfa romeo	chrysler	honda	norwich union	talk talk
alliance & leicester	citigroup	house of fraser	nowtv	tesco
allied irish bank	citroen	hsbc	npower	thameslink
amazon	coca-cola	hyundai	ofgem	the post office
american airlines	confused.com	ikea	oxfam	t-mobile
american express	co-op	jaguar land rover	paypal	topshop
anglian water	credit suisse	jet2	pc world	toyota
apple	currys	jp morgan	peugeot	tsb
argos	debenhams	kia	plusnet	tui
asda	deloitte	klm	porsche	uber
aston martin	deutsche bank	kpmg	powergen	unilever
audi	direct line	ladbrokes	premier league	united airlines
auto trader	dixons	land rover	prudential	uswitch
aviva	edf	legal & general	pwc	utility warehouse
axa sun life	eon	lehman brothers	quickquid	vanquis
bank of america	fox	littlewoods	renault	vauxhall
bank of cyprus uk	glaxosmithkline	liverpool victoria	rolex	viagogo
bank of ireland	the rac	lloyds bank	rolls-royce	virgin

bank of scotland	easyjet	marks & spencer	royal & sunalliance	vodafone
barclays	ebay	mastercard	royal bank of scotland	volkswagen
bbc	edf energy	mcdonald's	royal mail	volvo
bentley	equifax	mercedes-benz	ryanair	waitrose
bhs	ernst & young	merrill lynch	safeway	western union
big five	esso	met office	sainsbury's	wonga
big six	experian	metro bank	samsung	woolworths
bmw	financial times	microsoft	santander	yahoo
bp	flybe	moneysupermarket	savills	yodel
bradford & bingley	ford	morgan stanley	scottish power	yorkshire bank
british airways	gmtv	morrisons	scottish widows	zoopla
british energy	go compare	mothercare	shell	
british gas	goldman sachs	national grid	sky	
bt	google	national lottery	skype	
building societies	green flag	nationwide	south west trains	
bupa	gumtree	natwest	southeastern	

Table A2: List of the well-known firms used.

The organization entity nodes in the database that contained these firm names were given an additional property flagging them as a well-known firm. The entity extraction tools would identify firms mentions such as “Barclays”, “Barclays Bank” and “Barclays Bank plc” as separate

entities. To capture these variations on the same firm, if an organization entity name contained the letters “barclay” then that entity node would be classified as being a well-known firm. This would capture most of the variation of mentions on firm names. Firm names such as “eon” that may be contained as parts of other firm names were given a strict search criteria from manually checking the different variants of eon that appeared, e.g. “e.on” and using just those.

Another point about measuring well-known firms through the entity extraction tools is that if an article mentions “Barclays” three times, only one entity of “barclays” will be created for the article. Therefore, counting the number of times well-known firm entities are mentioned in an article only gives you the number of unique firm mentions in the article – or at least unique spelling variations of that firm name, as if “Barclays” and “Barclays plc” both appear in the article, this contributes towards two well-known firm mentions. So, the total well-known firm count in an article reflects to some extent the number of different firms mentioned within the article.

Each article node was given an additional property on how many well-known firm entities it contained. This number was summed over all rip-off articles to give the time series on the well-known firm count. As we saw in the results, the sum of well-known firm mentions over all rip-off articles peaks in 2018. This might indicate a wider focus on the above well-known firms in rip-off articles, rather than just mentioning the same handful of firms more frequently.