

Discourse about welfare benefits*

Simon Löwe

Burning Glass Institute, Bala Cynwyd, PA 19004, USA

Emily Silcock†

Department of Economics, Harvard University, Cambridge, MA 02138, USA

Abstract

We show policy discourse can substantially alter economic outcomes even when policy itself is unchanged, focusing on media discourse around welfare benefits. First, using 32.1 million newspaper articles from four countries, we characterise welfare discourse along three dimensions: framing, slant, and tropes. Discourse varies sharply across outlets: individual frames range from 1% to 33% of articles, and negative slant from 18% to 85%. 83% of negative stories invoke one of fourteen recurring welfare tropes. Second, we study the economic effects of this discourse, holding policy fundamentals fixed. Using an event-study design that exploits the timing of welfare coverage in *The Sun* newspaper, we find that Universal Credit applications fall by 7–8% after stories portraying individual claimants negatively, but not after policy or statistics frames, or positive stories. The response is stronger for stories that invoke a trope. Producing comparable changes through policy would require a 4–10% increase in benefit generosity.

Keywords: media stigmatization, media persuasion, policy discourse, framing, media slant, narratives, welfare stigma, benefit non-take-up.

JEL Codes: D83, H53, I38, L82, Z13

*This version: October 10, 2026.

†Corresponding author. Email: emilysilcock@fas.harvard.edu.

The world that we have to deal with politically is out of reach, out of sight, out of mind. It has to be explored, reported, and imagined.

— Walter Lippmann, *Public Opinion* (1922), p. 29

The information that reaches the public is never a full record of important events and developments in the world. It is, rather, a highly selective and stereotyped view of what has taken place. It could hardly be otherwise. But even if it could, the public would have little desire to be kept closely informed about the vast world beyond its personal experience. It requires news presentations that are short, simple, and highly thematic – in a word, stereotyped.

— John Zaller, *The Nature and Origins of Mass Opinion* (1992), p. 7

From conscription to vaccinations, Prohibition to Obamacare, the fundamentals of public policies are never communicated in a vacuum. Instead, they are debated, reported and imagined; what reaches the public is a politician making an argument about a policy, a newspaper editorialising on it, or a friend describing how it affects them. Unlike the fundamentals, this discourse may frame a policy in a particular way (Moreno-Medina et al., 2025), take a particular slant (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010), or tie into an oft-repeated narrative (Shiller, 2017, 2020).

Yet for many policies, while we know much about how changing policy fundamentals can change economic outcomes, we know little about the effects of changing the discourse around them.¹

In this paper, we focus on media discourse, as this remains the primary source of information on government policies.² Among these policies, we study media discourse around welfare benefits, as social protection is the largest area of government spending across OECD countries, on average,³ and a perennial topic of public debate. We begin by examining what

¹For example, we know that abortion bans have affected fertility and migration decisions (Dench et al., 2024, 2025), but how has the ‘pro-life’ and ‘pro-choice’ discourse affected individual decisions directly? We know that wealth taxes lead the rich to shift assets offshore (Alstadsæter et al., 2019; Londoño-Vélez and Ávila-Mahecha, 2025), but how have the prominent recent debates over taxing the wealthy shaped avoidance behaviour even where no such tax was ever enacted?

²In 2026, averaged across 11 policy topics, more than 40% of Americans said that their primary source of policy information was the news media, more than double the share who cited getting information directly from the government (17%), social media (14%) or friends, family and acquaintances (8%). Pew Research Center, *Americans say keeping up with key news topics is crucial. But fewer say they’re highly informed themselves*, 2026. <https://www.pewresearch.org/short-reads/2026/06/05/americans-say-keeping-up-with-key-news-topics-is-crucial-but-fewer-say-theyre-highly-informed-themselves/>.

³OECD, *Government at a Glance 2025*, “Government expenditure by function (COFOG)”. See https://www.oecd.org/en/publications/2025/06/government-at-a-glance-2025_70e14c6c/full-report/government-expenditure-by-function-cofog_d2b167d4.html.

this discourse looks like. Then we explore how variation in this discourse affects economic decisions while policy fundamentals are fixed.

In the first part of the paper, we characterise media discourse around welfare benefits along three dimensions: the frames it uses, its slant towards claimants, and the recurring tropes it draws on. We start with 32.1 million newspaper articles from sixteen of the highest-circulation newspapers from the US, UK, Germany and France, from 1995 to 2022. Using a custom natural language processing pipeline to classify articles, we find that welfare is a prominent policy topic; one in 57 articles mentions welfare benefits, around a third of the number that mention country leaders. The average newspaper has at least one article mentioning welfare benefits on 81% of days.

Narrowing down to articles that are mainly about welfare benefits, we split these into three frames: articles focusing on policy, articles focusing on statistics, and articles focusing on private individuals, a split motivated by the episodic versus thematic framing distinction of Iyengar (1991). We find newspapers within and across countries vary substantially in their use of these frames. An important divide is between tabloids, which are mass-market newspapers that tend to run short and sensational articles, and broadsheets, which tend to run longer and more neutral ones. In tabloids, 22% of articles use the individual frame, compared to only 7% in broadsheets. The UK has the highest share of individual frames and France the lowest.

Within the discourse on individuals, we look at slant towards people receiving welfare benefits (Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2010; DellaVigna and Kaplan, 2007).⁴ Tabloid newspapers contain markedly more negative slants, with 75% of articles being negative, compared to only 26% in broadsheets.

We then decompose negative articles into specific tropes. A trope is a recurring theme or character in storytelling. Previous work from outside economics has identified recurring characters in the welfare benefit discourse, such as the “welfare queen” (Levin, 2019). Further, Shiller (2017, 2020) argues that it is the recurrence of stories, rather than any one telling, that drives economic effects. To identify these recurring characters, we use an agentic clustering method (Löwe and Silcock, 2026). This uses a data-driven approach to create a taxonomy of tropes and then classifies articles into one of these tropes (or none). We find fourteen tropes, which account for 83% of negative articles about individuals. Examples include the *Prolific Parent* – parents with very large families living on benefits and portrayed as having children irresponsibly at the taxpayers’ expense (5.8% of articles) – and the *Con-*

⁴Slant is less explicit in policy and statistics frames, and we do not try to detect it.

cealed Death – individuals hiding the death of a relative to continue claiming their benefits (3.6% of articles).

In the second part of the paper, we ask how these frames, slants, and tropes affect real-world behaviour. We primarily focus on how this discourse affects individuals when they are deciding whether to take up welfare benefits. Non-take-up is a pervasive problem (Currie, 2006). In the US, non-take-up of means-tested welfare benefits ranges from roughly 16% for the Supplemental Nutrition Assistance Program (SNAP) to about 72% for cash assistance (TANF), and it exceeds 50% in Germany, Austria, and Finland and reaches 70–80% in Sweden (Ko and Moffitt, 2024). This is a large financial decision; a household in the US that decides to forgo SNAP loses \$258 a month,⁵ while a household in the UK that decides to forgo Universal Credit loses £660 a month (2019 values). Therefore we see this policy area as a demanding case for testing the effects of discourse. We also consider the impact on disability-related hate crimes and job search.

We focus on England and use data on Universal Credit, a composite means-tested benefit and the UK’s largest working-age welfare program. We obtained daily data on the number of applications to Universal Credit through Freedom of Information requests to the UK Department for Work and Pensions (DWP). For the discourse, we draw on articles from the British tabloid newspaper *The Sun*. For most of the period of study, this was the paper with the highest circulation in the UK, as well as the one with the most negative coverage of benefit recipients across all newspapers in our sample.

Using an event-study approach, we examine whether applications to Universal Credit are affected by discourse around welfare benefits. Our empirical strategy exploits the high-frequency variation in the timing of publication of newspaper articles. The publication of an article about welfare benefits is unlikely to be independent of long-run take-up trends; however, we argue that it is exogenous to the number of people applying for benefits on a given day, compared to the day before.

In our first set of results, we consider how Universal Credit applications respond to different frames in the discourse around welfare benefits. We find a large and significant decrease in the number of applications in the days following a story that uses an individual framing. On average, the number of applications decreases by 6–8% in the two days after publication, before returning close to the pre-publication baseline by the third day. By

⁵Mean monthly SNAP benefit across all participating households, fiscal year 2019. Cronquist, Kathryn (2020), *Characteristics of USDA Supplemental Nutrition Assistance Program Households: Fiscal Year 2019*, US Department of Agriculture, Food and Nutrition Service, p. xv. See <https://fns-prod.azureedge.us/sites/default/files/resource-files/Characteristics2019.pdf>.

contrast, we find no significant change in applications following a story with a policy or statistics frame, and the difference from individual framing is statistically significant.

Second, we examine the effects of the slant of the discourse. Relative to the day before publication, there is a drop of 7–8% in the number of Universal Credit applications in the two days following publication of a negatively slanted story about an individual. We do not find a response to stories about individuals that have a neutral or positive slant, although these are rare in *The Sun*. These results are stronger in geographical regions with higher *Sun* readership. The point estimates are also larger for articles closer to the front of the newspaper (and therefore more likely to be read), although this difference is not statistically significant. The results are robust to alternative specifications, to choices about event selection, and to dropping any single event.

Third, we compare negative stories that invoke an identified welfare trope with those that do not. The response is stronger for stories that fit into a trope. Whether a story invokes a particular trope is not randomly assigned, but these results are consistent with the hypothesis that discourse has stronger economic effects when it is repeated.

The prevalence of these tropes means that coverage offers a selective view of who claims benefits (Gilens, 1999). In the UK, individuals in articles are two times more likely to have been claiming for more than two years, six times more likely to be claiming fraudulently, and seven times more likely to have four or more children than claimants in administrative data. There is no reason to expect media to be representative (Armona et al., 2026); readers would likely be uninterested in articles that reported the typical claimant. Nonetheless, for fraud, the one attribute where survey evidence allows a comparison, the selection is correlated with public perceptions of welfare recipients. Survey respondents in 2007 put the share of false unemployment benefit claims at 37% (Geiger, 2018), compared to a fraud rate of 2.5% among Jobseeker’s Allowance claimants that year.⁶ This may suggest that, seeing the same types of claimant again and again, readers come to believe these examples are representative of the welfare claimant type (Rabin, 2002; Bordalo et al., 2016; Mastroiocco and Minale, 2018). These perceptions matter beyond claimants’ own decisions. The tropes, which range from many variants of benefit fraud to immigrants exploiting the welfare system, mirror the perceptions of deservingness and identity that shape voters’ support for redistribution (Stantcheva, 2021; Fisman et al., 2020; Alesina et al., 2018, 2023; Akesson et al., 2026; Bonomi et al., 2021), and may also shape how readers treat the groups they portray.

⁶DWP, *Fraud and error in the benefit system: financial year 2016 to 2017 estimates*, Table 5 (percentage of cases with a fraud overpayment, 2007/08). See <https://www.gov.uk/government/statistics/fraud-and-error-in-the-benefit-system-financial-year-2016-to-2017-estimates>.

We test this hypothesis using the largest trope in our data, the *Fake Disability Claimant*, a series of stories about individuals who fabricate disabilities in order to claim welfare. Repeated stories of disabled claimants faking their conditions may make readers more hostile towards disabled people, which we can measure through disability-related hate crime. As no official daily series exists, we construct one through Freedom of Information requests to every police force in England. The publication of *Fake Disability Claimant* articles is associated with a 16% increase in crimes on the day after publication, an effect that is statistically significant but imprecisely estimated. By contrast, we do not find a change in hate crimes in the days surrounding articles on other (or no) tropes, which portray other groups.

The validity of these results rests on the assumption that the publication of a story about welfare benefits is exogenous to the number of applications on a given day. First, we may be concerned about reverse causality. However, we do not find it plausible that newspapers publish negative articles about benefit seekers because they anticipate that there will be a short-run fluctuation in the number of applications on a given day, relative to the previous ones. Second, we may be concerned about high-frequency omitted variables that both increase the probability of a negative article being published and decrease the number of applications. For this reason, we flexibly control for day-of-week patterns in the applications data, using periods outside the windows around the publication of a story, and remove periods around public holidays from our analysis.

A natural concern is that negative articles affect applications not through the discourse itself, but through the specific content they report. As many of the articles in our sample discuss cases of benefit fraud, one possibility is that stories simply deter those who were planning to claim fraudulently. However, this cannot account for the magnitude of our results. Fraud made up only 2.8% of Universal Credit cases in 2015/16, so even if every would-be fraudulent applicant read the story and was deterred, this mechanism could explain less than half of the fall we observe. In practice it would explain much less, as most of this fraud arose after the initial application, through unreported changes in circumstances. A second possibility is that negative articles also convey information about the existence of benefits and eligibility for them, and a lack of such information is itself a known driver of non-take-up (Bhargava and Manoli, 2015). However, any such information effect would raise applications, working in the opposite direction to the decline we find.

The fall in applications could reflect individuals delaying their application or being deterred from applying altogether. We see applications return close to baseline by the third day, and no corresponding jump in applications within six days, indicating those who are

deterred from applying continue to be deterred for at least this long. Our high-frequency identification strategy does not allow us to see effects beyond the tight windows around events, so deterred individuals may apply later. However, much of non-take-up is driven by delaying rather than never claiming (Dynarski, 2007; Sunstein, 2013; Narayan, 2020), and even delaying for six days translates into a substantial financial loss for those affected.

A single negative story costs each affected household at least £130 (€148, \$166), and as much as £15,982 (€18,229, \$20,406) if it deters them from applying altogether (2019 values). The lower bound is 7.5–15.0% of monthly income for the median claimant household. Households forgo sums of this size in response to a single newspaper story, showing that discourse can matter even when the stakes are high.

Summing the effects of individual negative stories in *The Sun*, we account for around 4.7% of non-take-up over 2013–2022. Interventions addressing the most-studied drivers of non-take-up – a lack of information (Holford, 2015; Anders and Rafkin, 2026; Bhargava and Manoli, 2015) and the complexity of application procedures (Kleven and Kopczuk, 2011; Deshpande and Li, 2019; Finkelstein and Notowidigdo, 2019) – each reduce non-take-up by 5–9%. Coverage in a single newspaper is close to this range, pointing to discourse as a first-order driver of non-take-up.

Back-of-the-envelope calculations using estimates from the UK setting closest to ours suggest that a comparable change in take-up would require a 4–10% increase in benefit generosity, although these estimates do not identify a causal response to entitlement (Adam and Brewer, 2005). Estimates using strategies to identify causal effects, but from less comparable programs—unemployment insurance in the US and social assistance in Canada—imply increases of 8–12% (Anderson and Meyer, 1997; Whelan, 2010). These comparisons suggest that ordinary variation in the discourse around a policy can shift its outcomes as much as a substantial change in the policy fundamentals themselves.

A fall in take-up is not automatically a welfare loss. Negative coverage has been thought to encourage job search, nudging potential claimants towards work rather than benefits. We find the reverse, with online job search falling following negative stories, which may be evidence of a general avoidance of, or withdrawal from, domains or identities related to unemployment. Besley and Coate (1992) further suggest that deterring some applications could be socially beneficial, as such ‘ordeals’ improve targeting. If negative coverage works as an ordeal, the fall in take-up should be concentrated among the better-off. We find no evidence of this either; the response to negative stories is no larger in richer than in poorer areas. The timing of coverage compounds these effects. Tabloids carry more negative

coverage precisely when unemployment is high, the point in the cycle at which the Bailly–Chetty logic prescribes more generous insurance, not less (Chetty, 2006). Overall, we find no support for the channels through which negative discourse might raise social welfare.

Our results connect closely to the literature on the social determinants of take-up. Individuals are more likely to claim when others in their family or network claim (Bertrand et al., 2000; Dahl et al., 2014; Celhay et al., 2025), which may reflect information passing through networks or a fall in the stigma of claiming. The idea that stigma imposes a psychological cost on benefit receipt, and so depresses take-up, is long-standing (Goffman, 1963; Moffitt, 1983), and survey evidence attests to its presence (Morrison, 2019; Reeves and de Vries, 2016; Baumberg, 2016).⁷ Yet causal evidence is scarce, as surveys are hard to interpret causally (Celhay et al., 2024) and experiments study take-up of an experimental transfer in the lab rather than of an actual benefit (Friedrichsen et al., 2018). Field experiments have tested researcher-designed messages intended to reduce stigma, but find increases or decreases in take-up, depending on the specific message chosen (Bhargava and Manoli, 2015; McQuillan and Moore, 2026). We show that naturally-occurring stigmatising coverage reduces applications for the UK’s main means-tested benefit. Our design does not identify the effect of stigma alone – articles may deter claimants through channels other than shame – but it establishes that the discourse widely held to produce stigma is a quantitatively important driver of non-take-up.

Our paper additionally contributes to the literature on the economic and social effects of media (DellaVigna and La Ferrara, 2015). Existing designs commonly identify the effect of exposure to a particular outlet, medium, or program (DellaVigna and Kaplan, 2007; Enikolopov et al., 2011; Martin and Yurukoglu, 2017; Drago et al., 2014; Strömberg, 2004; Gentzkow, 2006; Jensen and Oster, 2009; Levy, 2021; Kearney and Levine, 2015; La Ferrara et al., 2012; Chong and La Ferrara, 2009), including *The Sun* itself (Foos and Bischof, 2022).⁸ This literature establishes that media matters, but often treats its content as a black box; it is less informative about which features of coverage drive the response. A second set of papers opens this ‘black box’, showing that the specific content of coverage matters, from the choice of a single word (Djourelouva, 2023) to specific linguistic structures (Moreno-Medina et al., 2025), but measures the effect on attitudes and beliefs rather than behaviour. We

⁷Similar surveys have been conducted in other contexts. In the US, Stuber and Kronebusch (2004) report that more than half of individuals surveyed agreed that “Many people on welfare do not want other people to know they are on welfare”, and between 35% and 45% that “People in this country on welfare are lazy”.

⁸Outcomes studied include hostility towards marginalised groups (DellaVigna et al., 2014; Ivandic et al., 2019; Ang, 2023; Esposito et al., 2023; Bursztyn et al., 2019; Müller and Schwarz, 2021, 2023) and labour-market behaviour (Murphy, 2020). We add results on disability-related hate crime and job search.

bring these two strands together, asking what in the reporting itself changes a decision with real financial stakes. Closest to our paper are studies linking the volume of coverage of a topic to a real-world outcome; Eisensee and Strömberg (2007) show that news coverage of a disaster raises the relief it receives, and Philippe and Ouss (2018) and Lim et al. (2015) that coverage of crime shortly before a trial lengthens the sentences judges impose. We unpack the black box further. Where these papers vary the volume of coverage of a single topic, we observe a high-frequency stream of millions of articles, classify their content, and identify which attributes of that discourse (its frame, its slant, and whether it belongs to a recurring trope) drive a consequential economic decision. This lets us move from showing that media matters to characterising what about media discourse matters.

The rest of the paper is organised as follows. Section 1 describes our data. Section 2 presents a set of stylized facts about the discourse around welfare benefits. Section 3 then considers the effect of this discourse, setting out our event-study methodology and results. Section 4 provides a discussion of these results, and section 5 concludes.

1 Data

Our analysis focuses on four countries: the UK, France, the US and Germany. The UK and the US are liberal welfare regimes, with tax-financed, largely means-tested safety nets in the Beveridge tradition, while France and Germany are conservative regimes built on Bismarckian contributory social insurance (Esping-Andersen, 1990). The four countries also span different media environments. Our sample contains both tabloids and broadsheets; the distinction was historically one of page size, but as many broadsheets adopted more compact formats it has come to serve as a shorthand for style, with tabloids running shorter and more sensational articles and broadsheets longer and generally more neutral ones. The UK and Germany combine strong public service broadcasting with a mass-market national tabloid press, and in both the widest-reaching titles in our sample are tabloids (appendix figure 17). France and the US have few national tabloids and more commercialised, politically differentiated television news markets, and their widest-reaching titles are broadsheets.

1.1 Newspaper articles

We use all articles from sixteen of the highest-circulation newspapers across these four countries. Appendix A lists the source and time period for each title (table 2) and describes how we extracted the articles.




Figure 1: Overview of newspaper article classification pipeline.

The size of our dataset precludes manually coding each article. Further, language used to describe welfare benefits differs across countries. A US reader may not interpret the phrase “a 50-year-old IT worker who’s just signed on” as relating to welfare benefits, whereas a UK reader would. Therefore we develop a custom natural language processing pipeline, which is summarised in figure 1, with more details in appendix B. First, each of the 32.1 million articles in our dataset is passed through a finetuned language model to classify whether the article mentions any benefits or the welfare system. Then we pass the articles that mention benefits through a second finetuned model to classify whether the article is mainly about benefits. We use GPT to separate articles mainly about welfare benefits into those that use a policy frame, articles that use a statistics frame, and articles that frame around individuals. For those that are mainly about individuals, we further use GPT to extract characteristics of the individual, such as their gender, number of children, and the number of years they have claimed benefits, giving us survey-like demographic data on the people the discourse portrays. We then use GPT to classify the slant of the article towards the individual as positive, negative, or neutral.

Finally, for negative articles about individuals, we follow the agentic clustering method of Löwe and Silcock (2026), which first creates a taxonomy of welfare tropes from the data.⁹ The method samples articles, proposes groupings, audits them against unseen articles and revises them, dispatching a separate large language model agent to each task, with an orchestrating agent deciding at each step which to call. Unlike standard clustering methods, it can be instructed, allowing us to discover clusters of recurring character tropes, rather than other ways in which the articles might be similar. While there is no single correct way to cluster a corpus, the method achieves state-of-the-art performance on standard text-clustering benchmarks, where the true grouping is known. The procedure results in fourteen distinct welfare tropes. We then assign articles to one of these tropes (or none) using a further GPT classifier. More details on this method are given in appendix B.5.

All models are carefully evaluated on hand-labelled data, and the performance of each model, as well as training details for the custom models, is reported in appendix B.

1.1.1 *The Sun* newspaper

For our causal analysis, we focus on *The Sun* newspaper, a daily tabloid, historically known for topless page-three models and sensationalist headlines of questionable veracity. *The Sun* was the highest-circulation newspaper in the UK for most of the period of our analysis.¹⁰ The paper’s parent company, News UK, is owned by Rupert Murdoch’s News Corp. A standard edition cost £0.40 (€0.47, \$0.63) in 2013 and £0.65 (€0.75, \$0.90) in 2019.

The Sun contains the most negative content about people receiving benefits of any newspaper in our dataset (appendix figure 24). Value-laden language about welfare benefits often appears on the front page. Figure 2 shows various examples.

We focus our analysis on the National Edition, which covers England and Wales. Separate regional editions, with different content, are published for Scotland, Northern Ireland, and the Republic of Ireland. We use articles from 2013 to 2022, and pass them through the natural language processing pipeline described above.

We further remove articles that focus on two ‘out-groups’: immigrants and Islamic extremists. In the former case, we saw a number of articles related to the possibility of Eastern European immigrants coming to the UK to seek benefits. In the latter, there were a number

⁹Michalopoulos and Xue (2021) similarly quantify recurring motifs across a large corpus and relate them to economic outcomes, though they classify against a pre-existing catalogue of motifs rather than deriving one from the data.

¹⁰Audit Bureau of Circulations data, reported by *Press Gazette*. See <https://pressgazette.co.uk/news/uk-national-newspaper-sales-slump-by-two-thirds-in-20-years-amid-digital-disruption/>.



Figure 2: Selection of *Sun* front pages.

of articles, mainly related to one individual, calling for the removal of benefits for individuals who had pledged allegiance to the Islamic State. In both cases, the articles focused on groups that *Sun* readers would generally not feel themselves to be members of, making the expected effect of articles unclear. Readers may still perceive negative slant towards all benefit seekers, or they may interpret such articles as meaning benefits should belong to the ‘in-group’, making them more likely to apply. We remove articles about television programs for the same reason; their subjects are known to readers as television personalities rather than as ordinary claimants.¹¹ We also do not include reader letters, which convey the opinion of a single reader rather than forming part of the media discourse we study. We use GPT to identify all these types of articles, and full details are provided in appendix B.

Overall, there are 786 days with at least one article focused on policy, 242 days with articles on statistics, and 713 days with articles on individuals (603 days with negative coverage of individuals and 145 days with positive or neutral coverage). Appendix B.4 presents a selection of articles from our dataset.

1.2 Universal Credit applications

Universal Credit in England provides an ideal context in which to study non-take-up, as it is a “catch-all” benefit which replaced all working-age means-tested welfare benefits. This allows us to study welfare benefits in general, rather than a specific benefit.

Universal Credit was introduced with the Welfare Reform Act of 2012. It was initially launched in areas of north-west England starting in April 2013, and then was rolled out

¹¹In particular, in 2014, Channel 4 released a controversial documentary called *Benefits Street*, which focused on the lives of several families on benefits. *The Sun* ran multiple stories about these individuals, about both their benefit claims and their wider lives.

progressively across the country. To start with, Universal Credit was only offered to a limited range of claimants, principally single, working-age people with no children who were seeking work. This was known as ‘live service’. Appendix figure 19 shows the roll-out across parliamentary constituencies over the period 2013–2016. In November 2014, ‘full service’ Universal Credit began to be rolled out, available to all working-age individuals. Live service was available across England by May 2016 and full service by December 2018.

Universal Credit replaced six means-tested benefits, known as ‘legacy benefits’: Jobseeker’s Allowance, an unemployment benefit;¹² Income Support, for individuals with a low income not receiving Jobseeker’s Allowance; Housing Benefit, support with housing costs; Employment and Support Allowance, for people with limited capability for work because of sickness or disability; Working Tax Credit, a top-up to low earned income; and Child Tax Credit, support with the costs of raising children.

During both the ‘live’ and ‘full’ service phases, people could only apply for Universal Credit if they were not receiving any legacy benefit, or if a change in their circumstances, such as moving to a different local authority, meant they would lose their legacy benefits unless they made a new claim. People who had no changes in circumstance continued to receive legacy benefits, until the final phase of the roll-out, ‘managed migration’, which started in July 2022 and is still underway. Hence, during our period of study, all Universal Credit claimants are either new applicants, or those receiving legacy benefits who had a change of circumstance. The overall number of households receiving either Universal Credit or legacy benefits remained stable from 2015 to early 2020, after which there was an uptick during the COVID-19 pandemic (see appendix figure 20a).

Eligibility for Universal Credit broadly matches the legacy benefits it replaced, so individuals receiving Universal Credit can be categorised according to which types of legacy benefit they would have received. Appendix figure 20b shows the number of households (or individuals) that receive Universal Credit or legacy benefits by type of benefit. As of May 2021, 62% of households receiving working-age, means-tested benefits were receiving Universal Credit, with notable differences across benefit types. 96% of individuals receiving unemployment benefits were receiving Universal Credit, whereas only 34% of those receiving disability-related benefits were receiving Universal Credit. This suggests that the turnover of individuals on unemployment benefits is far higher than that of disability benefits.

The DWP have not published take-up statistics of Universal Credit. However, they

¹²Only income-based Jobseeker’s Allowance has been merged into Universal Credit. Contribution-based Jobseeker’s Allowance has remained as a separate benefit. Unless specified, when referring to Jobseeker’s Allowance, we are referring to income-based.

have continued to publish data on take-up of legacy benefits, which implies non-take-up of between 10% and 44%, as shown in appendix figure 20c. Take-up of Universal Credit may be somewhat higher, as one of the key motivations for the switch to Universal Credit was that a simpler benefit system would increase take-up.¹³

We use administrative data on the daily number of applications to Universal Credit, covering every day from 22 April 2013 to 31 March 2022. It records the number of online applications on a given day, regardless of whether they were successful or not. The vast majority of applications are made online, as a claim can be made by telephone only if the applicant cannot apply online. This data was obtained via Freedom of Information requests made to the DWP.¹⁴ The data is broken down at the postcode area level (first two postcode characters), for the 99 postcode areas in England.¹⁵ The DWP withholds the count for any day with fewer than five applications, to protect confidentiality, and we set these withheld counts to zero. This affects 13% of postcode-area days, most of them weekends: over a third in 2013–2015, around a fifth in 2016–2018, and 2% or less from 2019. Summed across England, it leaves 55 days in our sample with zero recorded applications, all before 2016.

In appendix figure 21a, we show daily applications over the period. There is an upwards trend, due to the progressive roll-out of Universal Credit. Much of the variance of the data is caused by a regular weekly pattern in applications. In appendix figure 21b, we zoom into a three-week period in February and March 2016 to illustrate the weekly pattern. The number of daily applications decreases more or less monotonically across the week, with the highest numbers on Monday, and then substantially lower numbers during weekends, as shown in appendix figure 21c. In addition, we see lower numbers of applications at the end of each year, during the Christmas and New Year holiday period. We carefully account for this variation in our causal analysis.

1.3 Other data

We also use data on the demographics of Universal Credit and SNAP recipients, unemployment, regional income, *Sun* readership, disability-related hate crimes, and online job searches. Full details of the data sources and cleaning are given in appendix D.

¹³DWP, *Universal Credit: welfare that works*, Cm 7957, November 2010. See <https://www.gov.uk/government/publications/universal-credit-welfare-that-works>.

¹⁴See https://www.whatdotheyknow.com/request/daily_number_of_applications_for_67.

¹⁵Most of the major cities have their own postcode area (except London which is separated into 8) with populations ranging from around 30,000 to 2,000,000 per area.

2 Stylized facts about the discourse around welfare benefits

The following features of media discourse around welfare benefits can be found in the data:

Fact 1: Welfare benefits are a common feature of media discourse. On average, there are 3.8 articles mentioning welfare benefits per day per newspaper. By comparison there are 12.9 mentioning the country leader. The average newspaper has at least one article mentioning benefits on 81% of days. Tabloid newspapers have a smaller share of articles mentioning benefits than broadsheets (figure 3), likely due to a lower share of news content and a higher share of sport and culture content.




Figure 3: Share of articles mentioning welfare benefits and country leaders, by newspaper.

Notes: Bars show the share of a newspaper’s articles in each category. The welfare categories come from the classifiers described in appendix B. An article mentions the country leader if its headline or text names the leader in office on the day of publication, or their title and surname.

Turning to articles that are mainly about welfare benefits, tabloid and broadsheet newspapers have nearly identical levels of coverage over time, as shown in figure 4. Tabloid newspapers have more total articles, as these are shorter, and therefore have a higher count of articles that are mainly about welfare benefits than broadsheets. The one exception to this




Figure 4: Share of articles mainly about welfare benefits, by newspaper type.

Notes: Tabloid and broadsheet averages are computed by first calculating the share of articles mainly about welfare benefits for each newspaper, then averaging within each type.

pattern is the COVID-19 pandemic, when the share of broadsheet articles mainly about welfare benefits rose sharply. This spike is driven entirely by the United States (figure 22 in appendix E excludes US newspapers). Pandemic relief in the US was government-provided (e.g. Pandemic Unemployment Assistance), and therefore classified as welfare benefits, whereas in the other countries it was paid through employers (e.g. the UK’s furlough scheme), and therefore not classified as welfare benefits.

Fact 2: Tabloids and broadsheets use different frames in their discourse. Although tabloids and broadsheets have a similar share of articles mainly about welfare benefits, the framing of this discourse varies greatly between the two. As shown in panel (a) of figure 5, on average, 22% of articles mainly about welfare benefits in tabloid newspapers use an individual frame, compared to 7% in broadsheets. These shares are stable over time.

The breakdown by newspaper is reported in figure 23 in appendix E. While policy frames are the largest category for all newspapers, the share of individual frames varies from 33% (*The Sun*) to 1% (*Les Echos*). UK newspapers have higher shares of individual frames than those in other countries, but in every country tabloids have higher shares than broadsheets.

Fact 3: Three quarters of stories about individuals are negative in tabloid newspapers, compared to only a quarter in broadsheets. Focusing on articles that are only about individuals, we find that tabloid newspapers are strikingly more negative than

broadsheets. As shown in panel (b) of figure 5, on average, 75% of articles about individuals are negative in tabloid newspapers, compared to 26% in broadsheets.

The breakdown by newspaper is reported in figure 24 in appendix E. *The Telegraph* is the only broadsheet with more negative than positive or neutral articles, at 79% of articles about individuals. Across newspapers with at least 30 articles about individuals, the negative share ranges from 18% in *The New York Times* to 85% in *The Sun*.¹⁶

Tabloid newspapers are also more negative in their coverage of individuals during times of higher unemployment. Table 4 in appendix E shows a 1 percentage point rise in the unemployment rate is associated with a 2.5 percentage point increase in the share of individual stories that are negative. For broadsheets the coefficient is indistinguishable from zero.

Fact 4: 83% of negative articles about individuals fit into one of fourteen prevalent welfare tropes. A trope is a recurring theme, character type, or plot device in storytelling. Coverage of welfare benefits has been previously described as relying on a small set of such recurring tropes. Work in the UK has identified recurring characters, such as the “benefit tourist”, the “feckless father”, the “welfare cheat” and “benefit broods” (Jensen and Tyler, 2015; Morrison, 2019), while in the US the “welfare queen” has been mythologised




Figure 5: Frame and slant of articles about welfare benefits, by newspaper type.

Notes: Tabloid and broadsheet averages are computed by first calculating the percentage for each newspaper, then averaging across newspapers within each type. In panel (a), an article can be assigned to multiple frames, or none. Panel (b) excludes articles also about policy or statistics, and *Les Echos*, which has none left.

¹⁶This restriction excludes *Bild*, *Le Figaro*, *USA Today*, *Le Monde* and *Libération*.

Table 1: The fourteen welfare tropes discovered by the agentic clustering procedure, each with a representative article summary.

Trope	Example article summary
The Fake Disability Claimant	<i>“A claimant fraudulently received £17,329 in disability benefits while demonstrating full mobility by marching and drumming in a kilt, receiving a 56-day suspended sentence.”</i>
The Secret Worker	<i>“He claimed nearly £39,000 in benefits over seven years while working, funding two households, and got 250 hours’ community service instead of imprisonment.”</i>
The Hidden Wealth Fraudster	<i>“A fraudster living abroad claimed over £85,000 in social security over six years without disclosing significant savings, funded a lavish lifestyle and received prison sentence.”</i>
The Concealed Partner Fraud	<i>“Mother of five falsely claimed £69,366 in single-parent benefits while living with partner; secret surveillance uncovered the deceit, resulting in a suspended prison term.”</i>
The Identity Fraud Ring	<i>“A claimant used six false identities over 12 years to fraudulently obtain about €480,000 in social welfare payments, uncovered by facial recognition software.”</i>
The Prolific Parent	<i>“Criticises welfare recipients receiving large benefits and housing despite having many children, arguing such parents harm their offspring’s future and justify stricter benefit reforms.”</i>
The Criminal on Benefits	<i>“A compulsive shopper received concurrent jail terms for embezzling almost €100,000 from an employer and faking welfare claims, with partial suspension due to repayment.”</i>
The Toothless System	<i>“Ordered to repay £57,405 at £6.85 per week—161 years—after admitting five years of fraudulent benefit claims; given a six-month suspended sentence.”</i>
The Hardworking Taxpayer vs. The Scrounger	<i>“An elderly homeowner contrasts his street’s past of stable employment and pride with its current cycle of benefit dependence, anti-social behaviour and community decay.”</i>
The Concealed Death	<i>“He fraudulently claimed benefits for his deceased spouse over 14 years, receiving €258,000; he pleaded guilty, received suspended sentences and must pay €80,000 compensation.”</i>
The Lavish Lifestyle on Benefits	<i>“A mother of six receives £20,000 in benefits, hires a gardener, spends on supplements, makeup, diet aids, and avoids paid work to care for children.”</i>
The Immigrant Welfare Exploiter	<i>“A family linked to extremist facing terrorism charges departed after costly taxpayer-funded housing and legal battles, prompting calls for rule changes to avoid similar expenses.”</i>
The Fabricated Illness in Children	<i>“A mother received a four-year sentence after falsely claiming her children were disabled to obtain £350,000 in benefits, funding an extravagant lifestyle.”</i>
The System Insider Abuser	<i>“A tax credits adviser used insider knowledge to falsely claim single-parent status and low income, securing £65,000 in benefits; none repaid; sentencing deferred.”</i>

Notes: Full definitions of each trope, including the boundary notes supplied to the GPT classifier, are given in table S1 of the online supplement.




Figure 6: Prevalence of welfare tropes by country.

Notes: Each bar shows the share of negative articles about individuals (excluding articles also about policy or statistics) that are assigned to each trope, computed separately for the UK and the US.

(Levin, 2019).

Shiller (2017, 2020) models narratives as spreading epidemically, so that it is their recurrence and persistence, rather than any single telling, that drives economic behaviour. Therefore we extract these recurrent character tropes from our data, as described in section 1.1.

Applied to summaries of negative articles about individuals, this method recovers fourteen welfare tropes, listed in table 1 with a representative article summary for each. The tropes that we find line up with several of those proposed in the prior literature. Fraud is a common theme, with nine out of fourteen tropes involving illegitimate claims. The remaining five tropes cover receipt that is legal but presented as undeserved: large families living on welfare, conspicuous spending by individuals on welfare, immigrants receiving welfare, criminals who receive welfare, and ‘hardworking’ taxpayers who do not receive welfare.

83% of negative articles about individuals fit into one of these tropes, meaning that the vast majority are organised into a small number of recurring templates.

The prevalence of each trope in the UK and the US is shown in figure 6. We restrict this

comparison to these two countries as there are too few negative articles about individuals in France and Germany to estimate trope shares reliably. Many of the tropes related to benefit fraud are common across both countries. For example, *Fake Disability Claimant* is the most common trope in each (16.1% of UK articles and 14.8% in the US). Abuse by system insiders is a far more common trope in the US than in the UK, driven by several large stories (e.g. the Long Island Rail Road disability-pension scandal). Tropes related to deservingness are more prevalent in the UK. For example, 6.2% of articles in the UK focus on the *Prolific Parent* trope, but this is almost absent in the US.


Fact 5: Discourse around individuals receiving welfare benefits is not representative of claims. As we extract the characteristics of the individual at the centre of each article, we can compare the people portrayed in the discourse with the claiming population, on the same attributes that administrative data record. The prevalence of specific tropes leads certain attributes to be overrepresented in the discourse about individuals receiving welfare benefits, compared to their level among the claiming population.

Figure 7 shows the comparison between claims and articles for the UK and the US. (The number of articles about individuals in France and Germany is too low for comparison.) The black diamond shows the percentage of people claiming Universal Credit (in the UK) or SNAP (in the US) with a particular attribute. For example, 42% of people claiming Universal Credit in the UK have been claiming for more than two years.¹⁷ The red dot shows the percentage of negative articles about individuals receiving benefits where the individual has that attribute. We find that 93% of negative articles concern individuals who have been claiming for more than two years, around 2.2 times the level among claimants. The green dot shows the same but for positive and neutral articles.


Interestingly, we find that women and non-white individuals (in the UK) are underrepresented in the media discourse. While these groups have traditionally been recipients of elevated scrutiny in the media discourse around welfare benefits (Gilens, 1999), it seems this is no longer the case.

Instead the media discourse predominantly overrepresents mutable characteristics, such as the number of children, the length of claims, and claiming fraudulently. The prominence of fraud echoes Gavin (2021), who documents the scale of UK benefit-fraud coverage.

¹⁷For claim length in the UK we use the share implied by the spell-length model of appendix H.2 rather than the raw share of the caseload, which over our sample is 26%. Universal Credit was still rolling out for most of the period, so its caseload consisted largely of recent entrants and could not yet contain many claimants of two years’ standing.



(a) United Kingdom



(b) United States

Figure 7: Comparison of characteristics in newspaper articles about benefit claimants versus administrative statistics on welfare claimants in the UK and the US.

Notes: Points show the share of articles about individual claimants with each characteristic, separately for negative articles (red) and positive or neutral articles (green), against the share of claimants with that characteristic in administrative data (black diamonds). The figure beside each point is the ratio of the two. UK claims are from DWP Universal Credit and Fraud and Error statistics; US claims from SNAP Quality Control Public Use Files (FY1996–2022) and the 2022 Survey of Income and Program Participation. Lottery odds are 1 in 45,057,474 (UK Lotto) and 1 in 292,201,338 (US Powerball). Articles with unclear or unknown values are excluded, as are articles about fraud cases from the unemployed metric, since in those articles the individual's employment status is the substance of the allegation rather than an ordinary claimant characteristic. UK claim length is the share implied by the spell-length model of appendix H.2, since the raw caseload share is distorted by Universal Credit roll-out. Fraudulent claims, non-white and unemployed are shown for the UK only. US fraud statistics are not published, and we have no reliable race measure for US articles. The unemployed row excludes articles about fraud cases, which account for almost all negative US articles, leaving too few to report. Variable definitions and sources for both panels are given in appendix E.1.

There is no reason to expect that the media discourse should be representative of population characteristics. Stories about fraud, large families and lottery winners may be inherently more newsworthy than those about other individuals (Armona et al., 2026).

However, readers may extrapolate from salient examples to the population as a whole (Bordalo et al., 2016; Rabin, 2002), and a growing literature shows that misperceptions of who claims welfare, rather than material self-interest, shape attitudes towards redistribution (Kuziemko et al., 2015; Stantcheva, 2021; Alesina et al., 2023).

Consistent with this, for benefit fraud, the one attribute for which comparable survey evidence exists, British perceptions track media representation more closely than the actual level of claims. Individuals in articles are six times more likely to be claiming fraudulently than claimants in administrative data, and survey respondents in 2007 put the share of false unemployment benefit claims at 37% (Geiger, 2018), against a fraud rate of 2.5% among Jobseeker’s Allowance claimants that year.¹⁸

This is far from a causal analysis. Perceptions and coverage may both respond to the same underlying events, and newspapers may print what their readers already believe, but the pattern is consistent with readers believing these examples are representative of the welfare claimant type.

3 The effect of discourse on welfare benefit take-up

The previous section measured several features of the discourse about welfare benefits. We now test which of these features matter.

For this analysis we focus on the British tabloid *The Sun*, the most negative newspaper in our sample and the highest-circulation newspaper in the UK for most of our study period. We primarily study the effect on welfare benefit take-up, as this is an individual decision, and one that carries a heavy price tag. A household that chose not to apply for benefits in the UK would forgo £660 (€753, \$843) a month on average.¹⁹ We examine applications to Universal Credit, the UK’s catch-all working-age benefit. The general nature of this benefit

¹⁸Both administrative figures come from the DWP’s *Fraud and Error in the Benefit System* statistics. Each year the Department reviews a random sample of cases, checking each claim against government and local authority records and interviewing the claimant. Fraud is recorded where it is proven, or where the claimant ends or changes the claim after the review and there is strong reason to believe fraud was the cause. The 2007/08 Jobseeker’s Allowance rate is from *Fraud and error in the benefit system: financial year 2016 to 2017 estimates*, Table 5 (percentage of cases with a fraud overpayment). See <https://www.gov.uk/government/statistics/fraud-and-error-in-the-benefit-system-financial-year-2016-to-2017-estimates>.

¹⁹The mean award for households where Universal Credit is in payment, February 2019; see appendix H.1.

allows us to focus on the benefits system as a whole (see section 1.2).

3.1 Event-study design

We estimate the effect of stories on the number of applications for Universal Credit using an event study in the window around the date of publication of the first article in a story.

Specifically, we estimate regressions of the form:

$$\ln(\text{Applications}_t) = \alpha_{e(t)} + \sum_{\substack{\tau=-\underline{w} \\ \tau \neq -1}}^{\bar{w}-1} \beta_\tau D_{\tau,t} + X_t + \varepsilon_t \quad (1)$$

where t denotes calendar time and $e(t)$ the event whose window contains day t . An event is a window of days around the date of publication of a story, with $\underline{w}$ days before and $\bar{w}$ days from publication, including the publication day itself. Relative time within the window is denoted with τ , and $D_{\tau,t}$ is an indicator for day t lying τ days from its event. The reference day is the day before publication of the first article in a story. $\alpha_{e(t)}$ are fixed effects for events, and X_t collects controls for the calendar trend and day-of-week effects, described below.

The primary outcome variable, Applications_t , is the number of applications made to Universal Credit on day t . We only include applications made in England.²⁰ We exclude Scotland and Northern Ireland as they receive different editions of *The Sun* and we exclude Wales for consistency with the subsequent heterogeneity analysis for which we do not have data on Wales. We choose a logarithmic functional form to obtain relative effects, as the roll-out of Universal Credit across this period means that levels differ substantially across time.

We use Newey–West heteroskedasticity- and autocorrelation-consistent (HAC) standard errors with bandwidth 7. Appendix G.4 shows that inference is virtually unchanged under a data-driven bandwidth or under clustering by event.

Time trends As shown in appendix figure 21, applications to Universal Credit exhibit strong, but regular, weekly patterns, which may confound our results if negative articles are more likely to be published on particular days of the week.

Our preferred specification controls flexibly for the calendar trend and day-of-week effects: X_t collects a degree-five polynomial in calendar time and a set of day-of-week indicators, together with their interactions.²¹ We estimate this using all days in our sample, not just

²⁰One postcode area, Shrewsbury, contains parts of both England and Wales. We included this area.

²¹That is, $X_t = \sum_{j=1}^5 \gamma_j t^j + \sum_d \delta_d \mathbb{1}\{\text{dow}_t = d\} + \sum_{j=1}^5 \sum_d \phi_{jd} t^j \mathbb{1}\{\text{dow}_t = d\}$, with one day-of-week

those within the windows of events. Days outside every event window share a single event fixed effect, so they do not act as a comparison group for the β_τ ; they serve to identify X_t .

Alternatives for controlling for the time trend are considered in section 3.4, but yield nearly identical results.

Choice of window We define the event day as the date on which an article appears in *The Sun*. The ‘pre-period’ is the $\underline{w}$ days before it; the ‘post-period’ is the event day and the $\bar{w} - 1$ days after it.

These windows frequently overlap. We do not use any event for which another event, or the post-period of another event, lies within its pre-period (Morales, 2021). These overlaps threaten our identification of the effect of a given article when another article, or the post-period of another article, appears in its pre-period, as the supposedly untreated days are actually treated. In our main specification, however, we do use events which have another event, or the pre-period of another event, lying in their post-period. This is illustrated in panel (a) of figure 8, where the second event is used, as its pre-period is clear, but the third event is not used, as the second event window overlaps with its pre-period. We refer to this as a ‘one-sided’ overlap restriction. We allow events in the post-period, as these days remain treated, so the presence of another event does not threaten identification. Further, two events in quick succession are often related to the same story, with a lead article on the story on one day, followed by commentary on the same story on the following day. We particularly find this to be the case for bigger stories. Removing these events would dramatically reduce our sample size. It would also select the smallest – and least salient – stories, as it would select those that receive no follow-up article. The ‘one-sided’ window selection naturally selects the first article of a story for the treatment date. Thus our estimates can be properly interpreted as the effect of a story, or article cluster, rather than an individual article.

Nonetheless, our results are robust to a ‘two-sided’ window selection, where articles are only used if their event window does not overlap with that of any other article in the pre- or post-period. Results of this are shown in column (3) of table 9 in appendix G. Appendix table 5 accounts stage by stage for the reduction from the days with negative coverage of individuals to the events we use, of which the overlap restriction is by far the largest component.

A second choice is the length of the pre- and post-periods. As the length of the window determines which articles are kept and which are dropped, the choice of window length

category omitted. The estimation uses an orthogonal polynomial basis in place of the raw powers of t ; this spans the same space and leaves the β_τ unchanged.




Figure 8: Illustration of different ways of accounting for overlapping event windows. Green boxes represent events retained in the analysis and red boxes those removed.

presents a trade-off. If the windows are too short, we keep events whose pre-period is still affected by an earlier event. If the windows are too long, we remove events unnecessarily, reducing power.

For our main results, we use a pre-period of six days and a post-period of six days (including the event day), which balances these two concerns. We find that the effect of a story can be seen for two to three days, so this period is long enough to avoid pre-period contamination, while also avoiding unnecessary removal of events. Robustness to this choice is considered in appendix G.3.

Some of our results split the articles by an attribute of the article, such as whether it belongs to a trope. In these cases we subset the articles by the attribute first, and then apply the overlap and public holiday restrictions separately within each subset. A day on which articles of both types appear therefore counts as an event in both subsets. The exception is the split by page number, where each day is assigned to a single group using its frontmost article.

Identifying assumption This empirical strategy exploits the high-frequency nature of our data on Universal Credit applications. Our identifying assumption is that on average, the number of applications on the days following the first article in a story would have been the same as on the days prior to the article, if the article had not been published. This assumption is valid if the timing of negative articles is as good as random within a window of a few days (although this is not a necessary condition). That is, although the publication of negative articles about benefit seekers is unlikely to be exogenous to long-term trends in the number of applications, all that we require is that they are exogenous to changes in the

number of applications in the short window around publication. Our estimates are therefore the effect of one additional story, measured against the level of coverage prevailing at the time; our design does not identify the effect of that prevailing level itself. We do not claim to identify the effect of *The Sun* alone, but of the discourse more generally. Articles that are published by *The Sun* are sometimes carried by other newspapers, and may be re-reported on subsequent days, or may spread by social media or word-of-mouth.

There are three threats to this identification strategy. First, we may be concerned about reverse causality. However, we do not find it plausible that newspapers publish negative articles about benefit seekers because they anticipate that there will be a short-run fluctuation in the number of applications on a given day, relative to the previous ones.

Second, there may be measurement error concerns. It is possible that some of the articles appear on the *Sun* website the day before the print edition. Any such effect would fall on our reference day. As every coefficient is measured relative to that day, this would lower the baseline against which the fall in applications is measured, making our estimates a lower bound. An additional source of measurement error could come from the zero editing of our data on Universal Credit applications (section 1.2). From 2016, however, this is negligible: even if every withheld count were four applications, together they would make up less than 0.5% of national applications on the median day. In section 3.4 we show that our results are unchanged if we focus on the period from 2016 onwards.

Third, we may be concerned about high-frequency omitted variables that both increase the probability of a negative article being published and decrease the number of applications. For this reason, we control for day-of-week effects as discussed above. We exclude any event whose window contains a public holiday or a day between Christmas Eve and New Year’s Eve, as these affect benefit applications unpredictably and are hard to control for because there are too few of them to estimate their effects reliably. We also drop the event on 14 March 2020, whose post-publication window overlaps the early COVID-19 surge in Universal Credit applications. Our results are robust to the inclusion of this event, as shown in appendix G. Finally, we drop four events whose windows contain a day with zero recorded applications nationally, as the log of zero is undefined. A story is most likely to push a day’s applications to zero when its effect is large, so if anything this rule biases our estimates towards zero.

Other omitted variables may include big events such as elections, sports events or natural disasters. These are likely to decrease the amount of news space dedicated to coverage of social benefits. They may also decrease the number of applications for social benefits, if individuals are voting in elections, watching sports events, or recovering from natural

disasters. This would lead to a positive bias in results. Given that we find a negative effect of news articles on applications, this would bias our results towards zero.

3.2 Results

Discourse around individuals, policy and statistics We start by examining the application response to articles about individuals, policy and statistics. An article about an individual gives the reader a particular claimant to reason from, whereas one reporting a change to a benefit rate, or the size of the caseload, describes the system rather than the people in it.²² Iyengar (1991) argues that episodic framing (i.e. that focused on a specific individual or family) tends to make audiences attribute responsibility to the individuals depicted, unlike thematic framing (focused on trends, statistics or policy). Similarly, in lab experiments, Graeber et al. (2024) find that stories about individuals move beliefs further, and for longer, than the equivalent statistics. If readers respond to coverage in this way, applications should fall after articles about individuals, and not after articles about policy or statistics.²³

Articles about individuals are associated with a 7.7% drop in applications on the day after these articles, whereas we do not find a response to articles about policy or statistics




Figure 9: Event study of Universal Credit applications around *Sun* articles about welfare benefits. Left: articles about individuals. Right: articles about policy or statistics. The corresponding regression coefficients are reported in columns (1) and (2) of appendix table 6.

²²This distinction also echoes that made by Andre et al. (2026) who model stories with Directed Acyclic Graphs. Most of the policy or statistics discourse could be fitted into such a graph. By contrast, the discourse that focuses on individuals often makes no causal claim.

²³For the same reason, and as described in section 1.1.1, we exclude reader letters, articles about television programs and articles focused on two out-groups, immigrants and Islamic extremists, since readers are less likely to see themselves in the individuals these articles describe. Our estimates therefore capture the effect of negative coverage of claimants with whom *Sun* readers might plausibly identify.

(figure 9). The difference between the two sets of coefficients on the day of publication and the day after is statistically significant ($p = 0.015$). (Full regression coefficients are reported in columns (1) and (2) of appendix table 6.)

Negative, positive and neutral discourse We then turn to the slant of articles towards individuals receiving benefits. Among articles about individuals, we separate those with a negative slant from those with a positive or neutral one.²⁴ Given that we see articles about individuals depressing applications, we expect this effect to be concentrated among articles with a negative slant. Figure 10 shows applications on the days surrounding negative articles, as well as positive and neutral articles.




Figure 10: Event study of Universal Credit applications around *Sun* articles about individuals. Left: negative articles. Right: positive or neutral articles. The corresponding regression coefficients are reported in columns (3) and (4) of appendix table 6.

The pre-trends in the six days preceding the publication of the story with negative content about benefit seekers are flat and statistically insignificant, lending credence to our identifying assumption. Relative to the day before the publication of the story, there is a drop of 7–8% in the number of Universal Credit applications in the two days following the publication. On the third day, the number of applications has returned close to baseline. These decreases are both significant at the 1% level, with p -values of 0.005 and 0.001 respectively. These results imply that in the days following the publication of a negative story about individuals on welfare benefits, 7–8% of people who would have applied to Universal Credit are deterred from applying. We discuss the magnitude further and provide back-of-the-envelope computations about the implied social cost in section 4.2.

The equivalent event study for positive and neutral articles produces noisier results, due to the small number of such articles. However, we find no evidence of a response (either

²⁴Slant is rarely legible in policy and statistics articles, so we measure it only within the individual frame.

positive or negative) to the publication of these articles, although the difference between the two sets of coefficients is not statistically significant.

Heterogeneity by *Sun* readership Figure 11 also focuses on negative articles about individuals, but separates postcode areas with above-median *Sun* readership from those with below. (The corresponding regression coefficients are reported in columns (5)–(6) of appendix table 6.) The effect is larger in high-readership postcodes, both on the day of publication and on the day after, and a joint test rejects equality of the high- and low-readership coefficients on those two days ($p = 0.023$). The effect also lasts longer in high-readership postcodes, persisting for four days before returning to baseline, with the third and fourth days significant at the 10% level, whereas in low-readership postcodes it lasts for only two days. The larger effect in high-readership postcodes is consistent with more people there being exposed to the article directly, as well as more people encountering related discourse second-hand through conversation and social media. This difference in point estimates between high- and low-readership postcodes continues to hold when accounting for income (appendix figure 25).




Figure 11: Event study by *Sun* readership.

Notes: 1 of the 56 negative-individual events is excluded from this figure. Once postcode areas are split into subgroups, some subgroups have days with zero recorded Universal Credit applications; a zero count cannot be logged and is dropped from that panel. To keep both panels on an identical event set, we drop any event whose window contains a day with zero recorded applications in either panel.

Heterogeneity by article position in newspaper Similarly, articles are more likely to be read if they are closer to the front of the newspaper, especially if they are on the front page. Additionally, earlier articles tend to be on (what *The Sun* editors consider to be) more

sensational stories, which could be expected to be more negative towards benefit seekers. As figure 12 shows, we find that articles on page 5 or before are associated with a 19% drop in applications, whereas those on page 6 or after are associated with only a 6% drop, although the difference between the two sets of coefficients is not statistically significant.²⁵ The corresponding regression coefficients are reported in columns (7)–(8) of appendix table 6.




Figure 12: Event study by article position in the newspaper. Left: events where the earliest negative individual article appeared on pages 1–5. Right: events where the earliest such article appeared on page 6 or later.

3.3 Welfare tropes

We have seen that articles about individuals with a negative slant lead to a reduction in benefit take-up. Shiller (2017, 2020) argues that it is the repetition of narratives, rather than any single telling, that drives their economic effects. If so, we should expect a stronger response to articles that are part of a repeated trope than to those that are not. As we discussed in section 2, the majority of negative articles about individuals fit into one of fourteen repeated tropes. We therefore examine the differential response to trope and non-trope articles. Figure 13 shows the results.

We find a sharp two-day response (though only significant at the 10% level) to articles that are part of tropes. Results looking at articles that are not part of tropes are noisier, as most articles belong to a trope, but do not show any clear effects.

Whether an article is part of a trope is not randomly assigned – many other attributes of the article may correlate with this – so our results are not a conclusive test. Yet these results are consistent with the hypothesis that particular discourse is most powerful when it is repeated.

²⁵On days that have articles on more than one page, we use the earliest page with an article.

Taken together, these tropes portray claimants as undeserving, whether through fraud, idleness or lavish spending, or as outsiders to the community that funds the welfare system. These are precisely the perceptions that shape voters’ support for redistribution. Support falls when voters believe that the poor are poor through lack of effort rather than bad luck (Alesina et al., 2018; Stantcheva, 2021), that wealth is unearned (Fisman et al., 2020), or that beneficiaries belong to an out-group, such as immigrants or another racial group (Alesina et al., 2023; Akesson et al., 2026; Bonomi et al., 2021).²⁶




Figure 13: Event study of Universal Credit applications around negative *Sun* articles about individuals, split by trope membership. Left: articles fitting one of the fourteen tropes. Right: articles not fitting any trope. The corresponding regression coefficients are reported in columns (9)–(10) of appendix table 6.

Disability-related hate crimes As shown in section 2, 16.1% of negative articles about individuals fit the *Fake Disability Claimant* trope, making it the largest trope.

Here we look at whether articles in this trope are associated with disability-related hate crimes. Figure 14 shows disability-related hate crimes in the days following an article on this trope, compared to the days surrounding an article carrying another or no trope.

Some days have zero disability-related hate crimes, so we estimate these event studies using Poisson pseudo-maximum likelihood (PPML) rather than log-transforming the dependent variable. (Our main results are shown to be robust to this specification in appendix G.2.) The specification is otherwise that of equation (1). Because police forces enter the series at different points, its level trends upwards over the period, but the event fixed effects and time polynomial absorb this drift, so the effect is identified from variation within the short

²⁶Redistribution may in turn feed back into support for it. Avenancio-León et al. (2025) show that the racialised politicisation of the Food Stamp Program generated racial polarisation that persists half a century later, and Manacorda et al. (2011) find that a cash transfer raised support for the government providing it.

window around each article.

The results are noisy, but we do see a significant uptick in the number of recorded disability-related hate crimes on the day following the publication of a disability trope article. Such an article is associated with a 16% increase in expected hate-crime counts. By contrast, we do not find a change in hate crimes in the days surrounding articles not on this trope. The difference between the two sets of coefficients (for the day of publication and the day after) is significant, though only at the 10% level ($p = 0.057$). The corresponding coefficients are reported in columns (1)–(2) of appendix table 7.




Figure 14: Event study of disability-related hate crimes around negative *Sun* articles about individuals, split by trope. Left: articles carrying the *Fake Disability Claimant* trope. Right: articles carrying any other or no trope. The two panels share 2 events: dates that carry both a fake-disability article and an article with another or no trope.

3.4 Robustness

We perform a battery of tests, focusing on our results on negative articles about individuals. Full results are given in appendix G.

The publication-day and next-day declines remain significant at 5% under PPML estimation, and alternative ways of accounting for the day-of-week and time trends. This also holds when removing all events on Saturdays, which would be of most concern if we mis-specified the day-of-week trends, as the largest weekly drop in applications is between Friday and Saturday.

As discussed in section 3.1, we use a ‘one-sided’ removal of overlapping events: we remove an event if the event window of another event overlaps with its pre-period, but not its post-period. The declines remain significant under ‘two-sided’ removal, and under longer or shorter windows.

Leaving out any one event preserves significance at 5%, indicating that no single event drives the result. The same holds when retaining the event of 14 March 2020, which coincides with the early COVID-19 surge in Universal Credit applications, and is dropped in our main specification. No event in our sample falls in 2013, when Universal Credit was rolling out in north-west England only and application volumes were very low, so none of our estimates rests on that period. The declines also remain significant when subsetting to events from June 2016 onwards, once live service had been rolled out across all of England.

We also check alternative ways of computing standard errors, under which the declines remain significant. Given the low number of events, we additionally conduct inference with a Fisher Randomisation Test. We randomly draw events, following the day-of-week distribution we actually observe, and apply the full procedure as described above (including restrictions on dates due to window overlap and public holidays). The test yields p -values of 0.025 for the publication-day decline and 0.005 for the next-day decline.

4 Discussion

4.1 Delaying or deterring

We document a 7–8% reduction in the number of applications to Universal Credit over the two days after the publication of a negative story about benefit recipients, with the number of applications returning to baseline, or a little lower, by the third day. This is a change in the flow, not the stock, of applications and so, if there is no corresponding uptick in the number of applications, represents a permanent decrease in the number of benefit applicants. We see no such uptick in the six days following the publication of a story, suggesting that those who are deterred from applying are deterred for at least a week.

The short-term identification strategy we employ in this paper makes it impossible to see effects beyond the tight window of days around the publication of an article, so we do not know if individuals are simply delaying their application, or are fully deterred. However, this is not unusual with any determinant of non-take-up. Much of non-take-up involves delaying rather than never claiming, and procrastination has been shown to be a significant determinant of non-take-up (Dynarski, 2007; Sunstein, 2013; Narayan, 2020). Further, even if individuals are only deterred for a week, this corresponds to a significant financial loss, as we demonstrate in the next section.

4.2 Magnitude of effect

Cost to individuals Since benefits carry a directly observable monetary value, we can approximate the cost of these discourse effects with a back-of-the-envelope computation, without having to impute a value.

The average monthly Universal Credit award was £660 per household (€753, \$843), yielding a daily amount of around £22 (€25, \$28).²⁷ If we conservatively assume that a story only delays take-up by 6 days, the cost of a single stigmatising story to each affected household is £130 (€148, \$166). This is a lower bound: six days is simply how long we can observe the absence of a rebound in applications, and we see no evidence that households do in fact apply in the medium run. For comparison, this loss amounts to 7.5–15.0% of a month’s income for a household in the median income band reported in a DWP survey of recent Universal Credit claimants.

As an upper bound, if a story fully deters a household from claiming altogether, the cost rises to £15,982 (€18,229, \$20,406): the average monthly award multiplied by the 24.2 months a Universal Credit claim is expected to last (appendix H.2).

Effect on non-take-up Figures on non-take-up of Universal Credit are not published, but it may be somewhere between 10% and 44%, based on rates for the legacy benefits that it replaced (see appendix figure 20c). If we take the midpoint of this range, then around 73.0% of eligible people apply. Over the period, we have 750 ‘treated’ days (unique days with a negative individual article published on that day or the day before), which is 23.0% of total days. Since the published take-up figures average over treated and untreated days alike, they imply a counterfactual take-up rate, absent any negative story, of around 74.3% (appendix H.3). On a treated day, take-up is 5.5 percentage points below this, which is 20.4% of the baseline non-take-up of 27.0%. Scaling by the share of days that are treated, we explain around 4.7% of non-take-up overall.

This effect is comparable to those found for other drivers of non-take-up. Finkelstein and Notowidigdo (2019) reduce non-take-up of SNAP by 5% with an information intervention and by a further 7% with assistance in completing the application. In Bhargava and Manoli (2015), displaying benefit information reduces non-take-up of the Earned Income Tax Credit by 5% and simplifying the claiming worksheet by 9%. However, the effect that we document is a conservative estimate of the overall effect of the discourse around welfare benefits on

²⁷All Universal Credit figures in this section are 2019 values, which predate the COVID-19 pandemic. Appendix H gives the sources and details of each calculation.

non-take-up, as we are only looking at a single newspaper. Moreover, we compare days with and without one additional story, not a world with this discourse to one without it, and the effect of the discourse as a whole is likely to be larger than the sum of these small changes.

Comparison to effect of changes in policy fundamentals Reducing baseline non-take-up by a matching 4.7% would require a substantial change in policy fundamentals. In the UK setting closest to ours, Adam and Brewer (2005) imply that this would take an increase in entitlement of 3.9–6.0% for Family Credit or 6.2–10.0% for Working Families’ Tax Credit. Their estimates are associations rather than causal effects, and they note that earnings, and hence entitlement, may be endogenous to the take-up decision. Estimates using strategies to identify causal effects come from less comparable programs. Anderson and Meyer (1997) imply an increase of 9–12% would be needed in the weekly unemployment insurance benefit in the US, while Whelan (2010) implies an increase of C\$46–60 per month in Canada (1997 values; £20–26, €31–40, \$33–43), around 8–10% of the average entitlement in his sample, and Bruckmeier and Wiemers (2012) an increase of €37 per month in Germany (2006 values; £25, \$46). Both sets of estimates imply changes of a similar size. Appendix H.4 sets out each calculation.

4.3 Other content of stories

As mentioned above, the content of stories is not randomly assigned. Our identifying assumption allows us to compare days with articles to days without articles and identify the effect on benefit take-up. However, it does not allow us to compare days with articles containing tropes to days with articles not containing tropes, for example. Therefore it is possible that other attributes of articles, which are correlated with those that we have highlighted, are driving our effects. Here we consider the most likely candidates, and argue why none of them can account for what we find.

Information effect Articles about benefit seekers also contain information about benefits themselves. Therefore they may make people aware of benefits that they did not know about, or aware of their eligibility for benefits. Lack of information about benefits has been found to be a driver of non-take-up in previous work (Bhargava and Manoli, 2015). However, if an information effect is present, we would expect this to increase the number of applications after the publication of one of these stories, rather than decrease it. If anything, this means that our estimated effect is a lower bound on the effect of discourse.

Complexity effect Newspaper stories may also include information about the complexity of applying for benefits. This could include information on the process itself, or information on issues with the roll-out of Universal Credit, which could potentially increase individuals’ perceptions of the cost of applying, and therefore decrease their likelihood of applying for Universal Credit. Finkelstein and Notowidigdo (2019) and Deshpande and Li (2019) find that complexity costs can reduce benefit take-up. To rule out this mechanism, we looked for content which discussed the processes of application, or problems with the Universal Credit roll-out among our sample of negative articles. We found that none of the negative *Sun* articles in our sample touched on these topics at all. Therefore it seems unlikely that our effect is driven by an increase in perceived complexity.

Benefit fraud As mentioned earlier, many of the stories in our sample focus on cases of benefit fraud. Therefore another possible alternative mechanism for the effect of stories on application behaviour is that these stories deter would-be fraudulent applicants from applying. However, only 2.8% of Universal Credit claims were fraudulent in 2015/16. So even if every would-be fraudulent applicant saw the *Sun* story and was deterred by it, this mechanism could account for at most 43% of our observed effect. In practice it would account for much less. The vast majority of this fraud arose when claimants who were eligible at the point of applying later failed to report a change in their circumstances, such as a rise in earnings or a partner moving in, rather than from fraudulent applications. Only the latter could be deterred by a story at the point of applying.

4.4 Social welfare effects of discourse

Job search incentives Negative discourse about welfare benefits may have other socially desirable consequences, such as enforcing social norms that encourage (potential) claimants to exert job search effort (Lindbeck et al., 1999; Clark, 2003; Stutzer and Lalive, 2004). We examine whether negative *Sun* stories shift online job-search behaviour, using Google Trends searches within the job topic as a proxy for search intensity.²⁸

Figure 15 shows that in fact the opposite is true; individuals reduce their job search

²⁸As elsewhere, we use the log of the outcome. Because the log is undefined on the days when the Google Trends index is reported as zero (see appendix D), we drop events whose window contains such a day; this removes three of the negative-individual events and one of the positive/neutral events. Our Google Trends data cover all of 2013–2022, rather than the shorter period for which we have applications data, so more events fit inside the estimation window. The remaining event counts are reported in figure 15 and columns (3)–(4) of appendix table 7.

intensity in response to negative articles about individuals receiving welfare.

This finding is in line with McQuillan and Moore (2026) who find an increase in re-employment in a group encouraged to take up unemployment insurance. They argue that this effect is driven by these individuals being exposed to job-search requirements. This mechanism is unlikely to be driving the reduction in job search that we see here: while receipt may carry job search requirements, these do not come into force until the application is accepted, usually around a month later. Further, the magnitude of the effect suggests a reduction in search intensity not just by those who were deterred from applying by the story, but also by other job seekers, likely including those who already receive Universal Credit.

Instead, it is plausible that such a decrease in job search behaviour is driven by avoidance or withdrawal from domains and identities related to unemployment. Our findings indicate that, at least in the immediate aftermath of negative coverage of benefit recipients, withdrawal dominates any increased incentive to search in order to escape pressure from social norms. This effect may also drive part of that identified by McQuillan and Moore (2026).




Figure 15: Event study of Google Trends search volume in the “Job” topic around *Sun* stories about individuals. Left: negative stories. Right: positive/neutral stories. The corresponding regression coefficients are reported in columns (3)–(4) of appendix table 7.

In January 2019, 1.3 million households in England were receiving Universal Credit.²⁹ Averaging our two significant coefficients, we find that a negative story depresses job search by 4.9% for two days. If the claimants affected delay their job search by just these two days, and so remain on Universal Credit for two days longer than they otherwise would, a single story keeps households on Universal Credit for 123,959 additional days in total, at a cost to the public purse of £2.7 million (€3.1 million, \$3.4 million). Appendix H.5 sets out the assumptions behind this figure.

²⁹DWP, *Stat-Xplore*, Households on Universal Credit. See <https://stat-xplore.dwp.gov.uk/>. We use the same postcode areas as our applications data, so both cover the same geography.

Against this, the same story deters around 1,012 households from applying. As mentioned above, we do not know how long they are deterred from applying, but they would have to be deterred for at least four months, on average, to outweigh the cost from the reduction in job search effort.

Targeting effects of negative discourse It has been suggested that ordeals can improve the targeting of benefits to those who need them most (Besley and Coate, 1992). This hypothesis has received mixed empirical support, with some papers finding that ordeals deter those least in need of a public policy (Finkelstein and Notowidigdo, 2019; Shepard and Wagner, 2025), while others find the inverse (Bhargava and Manoli, 2015; Homonoff and Somerville, 2021; Deshpande and Li, 2019).

As we do not have individual characteristics of applicants, we cannot directly test how this hypothesis applies to the negative discourse about welfare policies. To shed some light on this issue, we assign each postcode area an average level of employment income for the 2018/19 tax year, constructed from constituency-level tax records as described in appendix D, and estimate our main specification separately in postcode areas above and below the median. Area income is only a coarse proxy for the need of marginal applicants, so this comparison is suggestive rather than a direct test of targeting.




Figure 16: Event study by 2018/19 average income, split at the median.

Notes: 3 of the 56 negative-individual events are excluded from this figure. Once postcode areas are split into subgroups, some subgroups have days with zero recorded Universal Credit applications; a zero count cannot be logged and is dropped from that panel. To keep both panels on an identical event set, we drop any event whose window contains a day with zero recorded applications in either panel.

Figure 16 shows no evidence of a larger effect of negative articles about individuals in

high-income areas than in low-income areas. This is also true when accounting for *Sun* readership (appendix figure 25). (As our outcome is in logs, the effect is not mechanically higher in low-income areas, where eligibility is higher.)

While we cannot know the individual income of those being deterred in these high-income and low-income areas, these results are consistent with low-income and high-income applicants being equally deterred by negative stories about welfare recipients, and do not support discourse assisting with the targeting of welfare benefits.

Unemployment rate As mentioned in section 2, tabloid newspapers contain more negative articles about individuals on benefits during times of higher unemployment. If the effects of negative articles on applications are homogeneous (an assumption we are underpowered to test), more people are deterred from applying during times of higher unemployment. This contrasts with the prescriptions of the Baily–Chetty formula (Chetty, 2006), which recommends more generous unemployment insurance when unemployment is high.

Overall, we find no support for the channels through which negative discourse about welfare benefits might raise social welfare.

5 Conclusion

Public policies are rarely communicated through their fundamentals; instead, they reach most people through discourse that frames them, makes certain attributes salient, and ties in to recurring narratives. In this paper we have studied this discourse for welfare benefits, one of the hardest test cases for whether such discourse matters, given the large sums at stake. Drawing on millions of newspaper articles from the UK, US, France and Germany, we first documented what this discourse looks like. Tabloid and broadsheet newspapers cover welfare at similar rates but in starkly different styles: tabloids far more often frame benefits through individual stories, which are also far more negative. Negative coverage concentrates into a small set of recurring tropes, which account for 83% of negative articles about individuals, and the claimants it depicts are more likely than those in administrative data to be long-term claimants, to be claiming fraudulently, and to have four or more children.

We then asked whether this discourse affects behaviour, using an event study around the publication of stories in *The Sun* and administrative data on daily Universal Credit applications. Stories framed around individual benefit recipients reduce applications by 7.7% on the day after publication, while stories framed around policy or statistics have no

detectable effect. Among stories about individuals, only those with a negative slant reduce applications, by 7–8% in the two days following publication, and more so for stories that fit an identified trope and in areas with higher *Sun* readership. Because non-take-up translates directly into forgone money, it also lets us price the effect: a single negative story costs each deterred household at least £130 (€148, \$166). Taken together, the results show that the discourse around a policy, and not only its fundamentals, shapes how it functions in practice.

Acknowledgements

We thank Luc Behaghel, Augustin Bergeron, Richard Blundell, Thomas Breda, Alex Bryson, Alice Butler-Warke, Melissa Dell, Xavier D’Haultfoeuille, Florian Foos, Ben Geiger, Philipp Ketz, James Morrison, Richard Murphy, Konstantin Poensgen, Pablo Querubin, Jesse Shapiro, Stefanie Stantcheva, Lesley Tassi, Shivram Viswanathan and David Yang for insightful comments and help with data access. We received excellent research assistance from Inioluwa Adeoluwa, Arno Golkar, Kartika Herscheid, Josh Kaplan, Daria Peters, François Praum-Gauvry, Abigail Summers and William Tarr. We also thank seminar participants at the Paris School of Economics and Harvard University. Support for this research was provided by the Chaire Travail and the Kenneth C. Griffin Economics Research Fund.

References

- S. Adam and M. Brewer. Take-up of Family Credit and Working Families’ Tax Credit: Final report. Technical report, HM Revenue and Customs, 2005.
- J. Akesson, R. W. Hahn, R. D. Metcalfe, and I. Rasooly. Race and redistribution in the United States: An experimental analysis. *Journal of Public Economics*, 255:105555, 2026.
- A. Alesina, S. Stantcheva, and E. Teso. Intergenerational mobility and preferences for redistribution. *American Economic Review*, 108(2):521–554, 2018.
- A. Alesina, A. Miano, and S. Stantcheva. Immigration and redistribution. *The Review of Economic Studies*, 90(1):1–39, 2023.
- A. Alstadsæter, N. Johannesen, and G. Zucman. Tax evasion and inequality. *American Economic Review*, 109(6):2073–2103, 2019.
- J. Anders and C. Rafkin. The welfare effects of eligibility expansions: Theory and evidence from SNAP. *American Economic Journal: Economic Policy*, 18(3):61–100, 2026.
- P. M. Anderson and B. D. Meyer. Unemployment insurance take-up rates and the after-tax value of benefits. *The Quarterly Journal of Economics*, 112(3):913–937, 1997.

- P. Andre, I. Haaland, C. Roth, M. Wiederholt, and J. Wohlfart. Narratives about the macroeconomy. *The Review of Economic Studies*, 2026. Forthcoming.
- D. Ang. *The Birth of a Nation*: Media and racial hate. *American Economic Review*, 113(6): 1424–1460, 2023.
- W. Antoun, F. Kulumba, R. Touchent, É. de la Clergerie, B. Sagot, and D. Seddah. CamemBERT 2.0: A smarter French language model aged to perfection, 2024. arXiv:2411.08868.
- L. Armona, M. Gentzkow, E. Kamenica, and J. M. Shapiro. What is newsworthy? theory and evidence. *American Economic Review: Insights*, 8(3):354–373, 2026.
- C. F. Avenancio-León, T. Howard, and W. Mullins. Dissecting racial politicization: Long-run evidence from the Food Stamp Program. Working Paper, Washington Center for Equitable Growth, 2025.
- B. Baumberg. The stigma of claiming benefits: a quantitative study. *Journal of Social Policy*, 45(2):181–199, 2016.
- M. Bertrand, E. F. P. Luttmer, and S. Mullainathan. Network effects and welfare cultures. *The Quarterly Journal of Economics*, 115(3):1019–1055, 2000.
- T. Besley and S. Coate. Understanding welfare stigma: Taxpayer resentment and statistical discrimination. *Journal of Public Economics*, 48(2):165–183, 1992.
- S. Bhargava and D. Manoli. Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11): 3489–3529, 2015.
- N. Boizard, H. Gisserot-Boukhlef, D. M. Alves, A. Martins, A. Hammal, C. Corro, C. Hudehot, E. Malherbe, E. Malaboef, F. Jourdan, G. Hautreux, J. Alves, K. E. Haddad, M. Faysse, M. Peyrard, N. M. Guerreiro, P. Fernandes, R. Rei, and P. Colombo. EuroBERT: Scaling multilingual encoders for European languages, 2025. arXiv:2503.05500.
- G. Bonomi, N. Gennaioli, and G. Tabellini. Identity, beliefs, and political conflict. *The Quarterly Journal of Economics*, 136(4):2371–2411, 2021.
- P. Bordalo, K. Coffman, N. Gennaioli, and A. Shleifer. Stereotypes. *The Quarterly Journal of Economics*, 131(4):1753–1794, 2016.
- K. Bruckmeier and J. Wiemers. A new targeting: A new take-up? *Empirical Economics*, 43(2):565–580, 2012.
- L. Bursztyn, G. Egorov, R. Enikolopov, and M. Petrova. Social media and xenophobia: Evidence from Russia. Working Paper 26567, National Bureau of Economic Research, 2019.
- P. A. Celhay, B. D. Meyer, and N. Mittag. What leads to measurement errors? evidence from reports of program participation in three surveys. *Journal of Econometrics*, 238(2):105581, 2024. Earlier version: NBER Working Paper No. 29184 (2021), “Errors in Reporting and Imputation of Government Benefits and Their Implications”.
- P. A. Celhay, B. D. Meyer, and N. Mittag. Stigma in welfare programs. *The Review of Economics and Statistics*, 2025. Forthcoming.
- R. Chetty. A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10–11):1879–1901, 2006.
- A. Chong and E. La Ferrara. Television and divorce: Evidence from Brazilian “Novelas”.

- Journal of the European Economic Association*, 7(2/3):458–468, 2009.
- A. E. Clark. Unemployment as a social norm: Psychological evidence from panel data. *Journal of Labor Economics*, 21(2):323–351, 2003.
- J. Currie. The take-up of social benefits. In A. J. Auerbach, D. Card, and J. M. Quigley, editors, *Public Policy and the Income Distribution*, pages 80–148. Russell Sage Foundation, New York, 2006.
- G. B. Dahl, A. R. Kostøl, and M. Mogstad. Family welfare cultures. *The Quarterly Journal of Economics*, 129(4):1711–1752, 2014.
- S. DellaVigna and E. Kaplan. The Fox News effect: Media bias and voting. *The Quarterly Journal of Economics*, 122(3):1187–1234, 2007.
- S. DellaVigna and E. La Ferrara. Economic and social impacts of the media. In S. P. Anderson, J. Waldfogel, and D. Strömberg, editors, *Handbook of Media Economics*, volume 1, pages 723–768. North-Holland, 2015.
- S. DellaVigna, R. Enikolopov, V. Mironova, M. Petrova, and E. Zhuravskaya. Cross-border media and nationalism: Evidence from Serbian radio in Croatia. *American Economic Journal: Applied Economics*, 6(3):103–132, 2014.
- D. L. Dench, M. Pineda-Torres, and C. Myers. The effects of post-Dobbs abortion bans on fertility. *Journal of Public Economics*, 234:105124, 2024.
- D. L. Dench, K. Lifchez, J. M. Lindo, and J. L. Liu. Are people fleeing states with abortion bans? Working Paper 33328, National Bureau of Economic Research, 2025.
- M. Deshpande and Y. Li. Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11(4):213–248, 2019.
- M. Djourelouva. Persuasion through slanted language: Evidence from the media coverage of immigration. *American Economic Review*, 113(3):800–835, 2023.
- F. Drago, T. Nannicini, and F. Sobbrío. Meet the press: How voters and politicians respond to newspaper entry and exit. *American Economic Journal: Applied Economics*, 6(3):159–188, 2014.
- S. Dynarski. The economics of student aid. NBER Reporter, No. 1, 2007.
- T. Eisensee and D. Strömberg. News droughts, news floods, and U.S. disaster relief. *The Quarterly Journal of Economics*, 122(2):693–728, 2007.
- R. Enikolopov, M. Petrova, and E. Zhuravskaya. Media and political persuasion: Evidence from Russia. *American Economic Review*, 101(7):3253–3285, 2011.
- G. Esping-Andersen. *The Three Worlds of Welfare Capitalism*. Princeton University Press, Princeton, NJ, 1990.
- E. Esposito, T. Rotesi, A. Saia, and M. Thoenig. Reconciliation narratives: *The Birth of a Nation* after the US Civil War. *American Economic Review*, 113(6):1461–1504, 2023.
- A. Finkelstein and M. J. Notowidigdo. Take-up and targeting: Experimental evidence from SNAP. *The Quarterly Journal of Economics*, 134(3):1505–1556, 2019.
- R. Fisman, K. Gladstone, I. Kuziemko, and S. Naidu. Do Americans want to tax wealth? evidence from online surveys. *Journal of Public Economics*, 188:104207, 2020.
- F. Foos and D. Bischof. Tabloid media campaigns and public opinion: Quasi-experimental evidence on Euroscepticism in England. *American Political Science Review*, 116(1):19–37,

2022.

- J. Friedrichsen, T. König, and R. Schmacker. Social image concerns and welfare take-up. *Journal of Public Economics*, 168:174–192, 2018.
- N. T. Gavin. Below the radar: A u.k. benefit fraud media coverage tsunami—impact, ideology, and society. *The British Journal of Sociology*, 72(3):707–724, 2021.
- B. B. Geiger. Benefit ‘myths’? the accuracy and inaccuracy of public beliefs about the benefits system. *Social Policy & Administration*, 52(5):998–1018, 2018.
- M. Gentzkow. Television and voter turnout. *The Quarterly Journal of Economics*, 121(3):931–972, 2006.
- M. Gentzkow and J. M. Shapiro. What drives media slant? evidence from U.S. daily newspapers. *Econometrica*, 78(1):35–71, 2010.
- M. Gilens. *Why Americans Hate Welfare: Race, Media, and the Politics of Antipoverty Policy*. University of Chicago Press, Chicago, 1999.
- E. Goffman. *Stigma: Notes on the Management of Spoiled Identity*. Prentice-Hall, Englewood Cliffs, NJ, 1963. Reprinted New York: Touchstone, 1986.
- T. Graeber, C. Roth, and F. Zimmermann. Stories, statistics, and memory. *The Quarterly Journal of Economics*, 139(4):2181–2225, 2024.
- A. Holford. Take-up of free school meals: Price effects and peer effects. *Economica*, 82(328):976–993, 2015.
- T. Homonoff and J. Somerville. Program recertification costs: Evidence from SNAP. *American Economic Journal: Economic Policy*, 13(4):271–298, 2021.
- R. Ivandic, T. Kirchmaier, and S. J. Machin. Jihadi attacks, media and local hate crime. CEP Discussion Paper 1615, Centre for Economic Performance, London School of Economics, 2019.
- S. Iyengar. *Is Anyone Responsible? How Television Frames Political Issues*. University of Chicago Press, Chicago, 1991.
- R. Jensen and E. Oster. The power of TV: Cable television and women’s status in India. *The Quarterly Journal of Economics*, 124(3):1057–1094, 2009.
- T. Jensen and I. Tyler. ‘benefits broods’: The cultural and political crafting of anti-welfare commonsense. *Critical Social Policy*, 35(4):470–491, 2015.
- M. S. Kearney and P. B. Levine. Media influences on social outcomes: The impact of MTV’s *16 and Pregnant* on teen childbearing. *American Economic Review*, 105(12):3597–3632, 2015.
- H. J. Kleven and W. Kopczuk. Transfer program complexity and the take-up of social benefits. *American Economic Journal: Economic Policy*, 3(1):54–90, 2011.
- W. Ko and R. Moffitt. Take-up of social benefits. In *Handbook of Labor, Human Resources and Population Economics*, pages 1–42. Springer, 2024.
- I. Kuziemko, M. I. Norton, E. Saez, and S. Stantcheva. How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508, 2015.
- E. La Ferrara, A. Chong, and S. Duryea. Soap operas and fertility: Evidence from Brazil. *American Economic Journal: Applied Economics*, 4(4):1–31, 2012.

- J. Levin. *The Queen: The Forgotten Life Behind an American Myth*. Little, Brown and Company, New York, 2019.
- R. Levy. Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3):831–870, 2021.
- C. S. H. Lim, J. M. Snyder, Jr., and D. Strömberg. The judge, the politician, and the press: Newspaper coverage and criminal sentencing across electoral systems. *American Economic Journal: Applied Economics*, 7(4):103–135, 2015.
- A. Lindbeck, S. Nyberg, and J. W. Weibull. Social norms and economic incentives in the welfare state. *The Quarterly Journal of Economics*, 114(1):1–35, 1999.
- W. Lippmann. *Public Opinion*. Harcourt, Brace and Company, New York, 1922.
- Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. RoBERTa: A robustly optimized BERT pretraining approach, 2019. arXiv:1907.11692.
- J. Londoño-Vélez and J. Ávila-Mahecha. Behavioural responses to wealth taxation: Evidence from Colombia. *The Review of Economic Studies*, 92(4):2624–2655, 2025.
- S. Löwe and E. Silcock. Agentic clustering: Controllable text taxonomies via multi-agent refinement, 2026. arXiv:2606.01255.
- M. Manacorda, E. Miguel, and A. Vigorito. Government transfers and political support. *American Economic Journal: Applied Economics*, 3(3):1–28, 2011.
- G. J. Martin and A. Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–2599, 2017.
- N. Mastrorocco and L. Minale. News media and crime perceptions: Evidence from a natural experiment. *Journal of Public Economics*, 165:230–255, 2018.
- C. McQuillan and B. Moore. Barriers to benefits: Unemployment insurance take-up and labor market effects. Job market paper, Princeton University and Stanford University, 2026.
- S. Michalopoulos and M. M. Xue. Folklore. *The Quarterly Journal of Economics*, 136(4):1993–2046, 2021.
- R. Moffitt. An economic model of welfare stigma. *American Economic Review*, 73(5):1023–1035, 1983.
- J. S. Morales. Legislating during war: Conflict and politics in Colombia. *Journal of Public Economics*, 193:104325, 2021.
- J. Moreno-Medina, A. Ouss, P. Bayer, and B. A. Ba. Officer-involved: The media language of police killings. *The Quarterly Journal of Economics*, 140(2):1525–1580, 2025.
- J. Morrison. *Scroungers: Moral Panics and Media Myths*. Zed Books, 2019.
- S. Mullainathan and A. Shleifer. The market for news. *American Economic Review*, 95(4):1031–1053, 2005.
- K. Müller and C. Schwarz. Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*, 19(4):2131–2167, 2021.
- K. Müller and C. Schwarz. From hashtag to hate crime: Twitter and antiminority sentiment. *American Economic Journal: Applied Economics*, 15(3):270–312, 2023.
- R. Murphy. Why unions survive: Understanding how unions overcome the free-rider problem.

- Journal of Labor Economics*, 38(4):1141–1188, 2020.
- A. Narayan. Does simplifying the college financial aid process matter? *Economics of Education Review*, 75:101959, 2020.
- A. Philippe and A. Ouss. “no hatred or malice, fear or affection”: Media and sentencing. *Journal of Political Economy*, 126(5):2134–2178, 2018.
- M. Rabin. Inference by believers in the law of small numbers. *The Quarterly Journal of Economics*, 117(3):775–816, 2002.
- A. Reeves and R. de Vries. Does media coverage influence public attitudes towards welfare recipients? the impact of the 2011 English riots. *The British Journal of Sociology*, 67(2): 281–306, 2016.
- M. Shepard and M. Wagner. Do ordeals work for selection markets? evidence from health insurance auto-enrollment. *American Economic Review*, 115(3):772–822, 2025.
- R. J. Shiller. Narrative economics. *American Economic Review*, 107(4):967–1004, 2017.
- R. J. Shiller. *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton University Press, 2020.
- E. Silcock and S. Löwe. Active learning for finding a needle in a haystack, 2026. Preprint.
- S. Stantcheva. Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369, 2021.
- D. Strömberg. Radio’s impact on public spending. *The Quarterly Journal of Economics*, 119(1):189–221, 2004.
- J. Stuber and K. Kronebusch. Stigma and other determinants of participation in TANF and Medicaid. *Journal of Policy Analysis and Management*, 23(3):509–530, 2004.
- A. Stutzer and R. Lalive. The role of social work norms in job searching and subjective well-being. *Journal of the European Economic Association*, 2(4):696–719, 2004.
- C. R. Sunstein. Nudges.gov: Behavioral economics and regulation. SSRN Working Paper 2220022, 2013.
- S. Whelan. The take-up of means-tested income support. *Empirical Economics*, 39(3): 847–875, 2010.
- J. R. Zaller. *The Nature and Origins of Mass Opinion*. Cambridge University Press, Cambridge, 1992.