White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of York.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/247000/

Version: Accepted Version

Article:
Meers, Jed, Stojanovic, Iva and Tomlinson, Joe (Accepted: 2026) Of Burdens and Bots: How Benefit Claimants Use Generative AI Chatbots. Social Policy & Administration. ISSN: 1467-9515 (In Press)

Reuse
This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: https://creativecommons.org/licenses/

Takedown
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

UNIVERSITY OF LEEDS
University of Sheffield
UNIVERSITY of York

Research Article **Social Policy & Administration**

Of Burdens and Bots: How Benefit Claimants Use Generative AI Chatbots

Dr Jed Meers, Iva Stojanović, and Professor Joe Tomlinson
University of York, UK, Harvard University, USA, King's College London, UK.

We know an increasing amount about the use of AI tools by officials working across the design and delivery of the welfare state, but far less is understood about how those accessing support are making use of this technology themselves. This article presents an analysis of 271 chats that current recipients of Universal Credit and Personal Independence Payment in the UK held with commercial chatbots, such as ChatGPT, Claude, and Gemini. It is widely recognised that maintaining a claim for social security support imposes a range of administrative burdens on claimants, and our 'chat mining' data show that claimants use AI tools in a way which both reflects and reduces the learning, compliance, and psychological costs of these burdens. The article argues that claimant use of these tools is likely to have significant and unpredictable implications for social security administration.

Keywords: Generative AI, welfare advice, Universal Credit, digital welfare state, administrative burden.

Introduction

Over a third of the UK public have used a generative AI chatbot as part of their 'day-to-day life' in the past month (Department for Science, Innovation and Technology, 2026). In common with most forms of digital services, a small number of platforms (particularly ChatGPT, Claude, Gemini, Copilot, and Grok) dominate the market, providing free-at-the-point-of-use tools that are widely used for practical guidance, seeking information, and completing writing tasks (see OpenAI, 2026). As these tools proliferate, so too does research examining their use across public services. A growing body of literature explores how generative AI tools are used by civil servants (see Bright et al., 2025; Larsen & Følstad, 2024), how the public interact with bespoke AI chatbots built into public services (see Følstad & Bjerkreim-Hanssen, 2024; Verne et al., 2022), and how their existence can further worsen the inequalities that result from digital exclusion (see Alexopoulou, 2025, pp. 145–146; Capraro et al., 2024; Wang et al., 2025).

Within this evolving body of research, this article examines a question that remains heavily underexplored: how do benefit claimants use publicly available generative AI chatbots in the course of their claim? This question is important for at least two reasons. First, most current research is from the point of view of officials. It examines how they incorporate (or do not

incorporate) generative AI in the design or delivery of public services, or, more informally (and possibly without the blessing of senior staff), use publicly available generative AI tools in their day-to-day work (for an evaluation of the latter in the UK, see Department for Work and Pensions, 2026c). Second, where research does explore claimant uses of generative AI, this tends to focus on bespoke tools built into public services, such as chatbots incorporated into existing government platforms (see Følstad & Bjerkreim-Hanssen, 2024; Verne et al., 2022). These applications of generative AI are, by definition, designed and controlled by public services themselves. This same control cannot be exerted over the comparatively unstructured publicly available generative AI chatbots such as ChatGPT, which can be used by any member of the public with the means to do so, well beyond the oversight of officials or whatever guardrails they prefer to have in place. This development promises to be particularly significant for research on social policy and public administration. A substantial literature charts the digitalisation of the welfare state, in which governments and their commercial contractors have digitalised the infrastructure of claiming, assessment, and payment (see Alexopoulou, 2025; Schou & Pors, 2019). Generative AI marks the latest stage in that trajectory, but its most widely used applications are commercial chatbots built by a small number of technology firms rather than by public bodies themselves, and they are already in the hands of claimants (or at least those with the capability to access them). As this paper goes on to argue, these tools are already reshaping the processes and day-to-day realities of citizen-state interactions that this literature explores. Where digitalisation to date has largely kept the design of these interactions within the state's own imperfect, yet predictable control, generative AI introduces a new intermediary that sits wholly outside it. This intermediary holds no formal accountability to claimants or the state for the guidance it gives, and its capabilities shift with every model release. Whether this ultimately helps claimants exercise their rights more effectively or compounds the inequalities already documented in this literature is a question this article returns to.

Drawing on a method designed to capture chat prompts ethically and securely (an approach we characterise as ‘chat mining’), we examine how benefit claimants are using these tools by analysing a sample of 271 chat exchanges shared by 181 Universal Credit (UC) and/or Personal Independence Payment (PIP) recipients in the UK. Our data suggest that benefit claimants’ use of these tools goes far beyond an alternative to web search or a simple means of seeking advice. Instead, what emerges is a wide-ranging set of applications that at once reflect and seek to mitigate the administrative burdens faced by people accessing public services and those identified by researchers in the social security context specifically: *learning* costs (finding information about entitlement, rules and social security programmes), *compliance* costs (demonstrating conditions of entitlement have been met, filling out forms and so on) and *psychological* costs (meeting the emotional toll of interacting with public services) (see Herd & Moynihan, 2018; Bennett et al., 2026). We argue that their use, currently out of sight from officials working in the welfare state and across public services more generally, has significant implications for the design of front-line services and the experiences of claimants themselves. In doing so, we extend the administrative burden framework by showing that commercial generative AI tools are taking a place alongside other kinds of support for meeting the costs of burdens, such as friends, family members, and advisers on whom claimants have long relied.

We develop this argument in four sections. The first situates this study in the wider literature on the use of generative AI tools in public services and the social security system specifically. We demonstrate that the burgeoning literature on the use of generative AI has to date neglected

claimants' independent use of widely available tools such as ChatGPT, Claude, and Gemini. Second, we set out the three types of administrative burden costs that we return to in the analysis of the chat prompts. Third, we detail our 'chat mining' method, a novel approach to capturing real-world claimant-AI interactions. Fourth, we turn to the findings, providing illustrative examples of chat prompts from participants in the study throughout to support our claim that they both reflect and seek to offset the costs imposed on claimants through administrative processes. A concluding discussion then draws out what these findings mean for how administrative burden is studied and produced, and argues that changes to the welfare state can come not only from the state's adoption of AI, but from how claimants are already using it themselves.

Generative AI use in public services

Although this article focuses on one specific kind of use of generative AI (claimants' use of publicly available, privately developed chatbots), a growing literature has begun to examine the broader role played by AI chatbots in welfare administration (e.g. see Agostino et al., 2022; Pors & Pallesen, 2021). Existing research has primarily developed along two lines: the design and implementation of chatbot assistants by public agencies themselves, and the use of AI systems within the day-to-day operations of state bureaucracies.

On the first strand, governments are increasingly designing and delivering their own AI-assisted chatbots. A 2022 UN e-Government Survey found that 69 countries had integrated AI chatbots into their national web portals (United Nations, 2022), and in 2024, 94 countries explicitly referenced AI in their national e-government strategies, up from 65 in 2020 (United Nations, 2024). Existing research has raised concerns about how chatbot systems may shape trust in public decision-making (Aoki, 2020), mediate experiences of welfare stigma (Considine et al., 2022), deepen forms of technological exclusion (Alexopoulou, 2025; Schou & Pors, 2019), and expand state surveillance capacities (Jørgensen, 2023; Rachovitsa & Johann, 2022; Kaun & Männiste, 2025). However, research also identifies the potential of chatbots to reduce administrative burden in government-citizen communication (Androutsopoulou et al., 2019) and of generative AI to bridge access-to-justice gaps caused by a lack of access to advice and support (Chien & Kim, 2024; F. Ryan & Hardie, 2024).

Most empirical work in this space examines chatbot implementations in specific government services, such as local government deployments in Latvia, Vienna and Bonn (van Noordt & Misuraca, 2019), municipal chatbots in Scandinavia (Følstad & Bjerkreim-Hanssen, 2024; Verne et al., 2022), and national portal systems in Estonia (Alishani & Homburg, 2025). A major systematic review of e-government chatbots found that their predominant use remains confined to information search, document retrieval, and basic service guidance (Cortés-Cediel et al., 2023), while broader reviews have documented chatbot adoption across local governments internationally, identifying information provision, consultation, and complaint-handling as the most common functions (Senadheera et al., 2025).

The second strand of this literature looks instead at how officials themselves use generative AI tools in day-to-day administrative and legal work. In the UK, government guidance has increasingly encouraged cautious experimentation with generative AI across public administration, while departments such as the Department for Work and Pensions (DWP) have developed internal policies governing staff use of approved and public-facing AI tools (Department for Work and Pensions, 2026a; UK Government, 2024). Much of this literature

frames generative AI as a tool capable of augmenting bureaucratic work by increasing efficiency, reducing administrative burdens, or improving access to complex systems. Choi et al. (2025), for example, find that Korean civil servants associate generative AI use with improved work performance and greater administrative efficiency. Similarly, a growing legal literature has explored the potential of generative AI to support drafting and translation tasks and assist legal professionals navigating resource constraints (Chien & Kim, 2024).

A smaller but emerging body of largely legal scholarship has begun to consider how citizens themselves may use publicly available generative AI systems when navigating legal and administrative problems. The closest to our approach is Kieslich et al.'s (2026) study of imagined uses of generative AI in legal conflict resolution. Drawing on scenario-writing exercises with citizens and legal professionals in the EU and US, they show participants repeatedly envisioning people turning to tools such as ChatGPT to draft complaints, interpret legal problems, seek procedural guidance, or replace legal advice they could not otherwise afford. However, these accounts do not specifically relate to welfare provision and focus on hypothetical situations rather than 'real world' uses of these tools. More broadly, existing legal scholarship has focused primarily on the risks associated with reliance on open-access generative AI tools, including inaccurate legal guidance, procedural mistakes, and the possibility that these technologies may deepen existing inequalities in access to justice (F. Ryan & Hardie, 2024). Yet despite growing concern about citizen reliance on commercial AI systems when navigating complex institutions, little work has examined how welfare claimants themselves are already using these tools in the course of making and maintaining benefit claims.

What emerges therefore is a literature that remains primarily concerned with chatbot systems embedded within public services themselves and the use of generative AI tools (and AI more generally) by civil servants. Much less is known about how claimants independently use publicly available generative AI tools such as ChatGPT, Claude or Gemini when navigating the welfare state, and whether these tools mediate the administrative burdens associated with claiming support. Given the widespread (and increasing) use of these tools across the population, this is a significant gap in our understanding that is likely to become even more acute over the coming years.

Making and Maintaining a Claim: Administrative Burdens

Having explored the key strands of research on generative AI tools in public services, we turn now to the literature on administrative burdens, which provides the analytical framework for the discussion that follows. Making and sustaining a claim for welfare benefits imposes a wide range of administrative burdens on claimants. These are explored in a field of its own, developed initially in the US context but now applied across welfare states internationally (see Bennett et al., 2026; Peeters, 2020). Although the subject of a renewed focus in the field of public administration in recent years, the idea has a much longer intellectual pedigree (for an overview, see Somers-Joce & Tomlinson, 2026, pp. 512–513). We draw on Herd and Moynihan's influential threefold typology of the costs associated with these burdens: learning costs, compliance costs and psychological costs (Herd & Moynihan, 2018; Moynihan et al., 2015). Before setting out each of the three costs in turn, we first introduce the two UK benefits on which this study focuses, since the character of these burdens in claimants' lives is shaped by the design of the benefits themselves.

Universal Credit (UC) is the UK's principal means-tested, working-age social assistance payment, claimed by over eight million people in and out of work, and has progressively replaced six 'legacy' benefits since its introduction in 2013 (Department for Work and Pensions, 2026b). Personal Independence Payment (PIP) is a needs-tested benefit paid to almost four million people towards the additional costs of disability, payable whether or not the claimant is in work, with entitlement determined by assessing the functional impact of a claimant's conditions against a set of 'descriptors' (Friskney et al., 2026). The distinction matters for what follows. The UK has no substantial social insurance tier, so UC operates as the final safety net for those who rely on it and the stakes of getting a claim right are correspondingly high (Bennett et al., 2026), while PIP's assessment-driven design places a premium on claimants' ability to articulate the impact of their conditions in the terms the system recognises (Day & Shaw, 2022).

Turning to the costs themselves, the first - learning costs - arise from the burdens of acquiring the information necessary to access and maintain entitlement to a public benefit. They are particularly significant in the social security context, where benefits regimes are governed by complex and frequently changing rules concerning entitlement, evidence requirements and the conditions attached to the ongoing receipt of support (Harris, 2013). In the UK, although UC was introduced as a simplification of the welfare system, in practice it remains complex and imposes considerable learning costs on claimants (Bennett et al., 2026). PIP claimants face their own learning costs, such as the distinction between the daily living and mobility components, the descriptors used to determine entitlement, and the rules around reviews and reassessments (Friskney et al., 2026). Importantly, claimants are often trying to navigate multiple sources of support at once across a range of government services, including at the local level, a problem Bennett (2026) describes as a 'multi-level maze'.

Second, compliance costs are the burdens of demonstrating that entitlement criteria have been met and of meeting the obligations attached to receiving support. They encompass form-filling, providing documentary evidence, attending appointments, and so on. The form is the archetypal site of compliance costs, with Brown et al.'s (2021) work drawing attention to how form design can impose substantial burdens of its own on claimants (see also Meers, 2024; Ryan, 2023). However, compliance costs extend well beyond the initial process of applying. In the UC context, they include the obligation to log job searches against a Claimant Commitment, to report changes of circumstances, and to upload evidence in the formats required by the digital system (Meers et al., 2025). In digital systems of this kind, this work does not disappear so much as shift onto claimants themselves, who become what Madsen et al. (2022) term 'accidental caseworkers'. For PIP claimants, compliance costs include the completion of the PIP2 questionnaire and the assessment process, where claimants are expected to articulate the impact of their condition against the relevant descriptors (Day & Shaw, 2022). Compliance burdens are therefore both episodic (taking place at moments of application, assessment and review) and a continuous part of the day-to-day experience of sustaining a claim.

Finally, psychological costs flow from the burdens of stress, frustration, stigma, and loss of autonomy interacting with public services can cause (Bennett et al., 2026, p. 37). They are perhaps the least visible of the three, but the literature on welfare claiming suggests they may be the most consequential, with the emotional toll of claiming shaping both take-up rates and the experience of claiming itself (Baumberg, 2016; Patrick, 2017). The psychological costs

associated with UC and PIP are well-documented in the UK: some claimants have described the sustained stress of waiting for a decision, the anxiety attached to assessment and reassessment, the fear of sanctions, and the stigma associated with aspects of claiming (Cheetham et al., 2019; Dwyer et al., 2020).

These three burdens interact and compound one another (Bennett et al., 2026, p. 29). For instance, a claimant who does not understand what evidence is required (a learning cost) is less able to meet the requirements of an application or review (a compliance cost), and may experience heightened anxiety or feelings of stigma as a result (a psychological cost). Likewise, these costs are not evenly distributed. Claimants with health conditions, disabilities, caring responsibilities, or insecure employment will often face the heaviest burdens at precisely the moments when they have the least capacity to absorb them (Heinrich, 2016). As we demonstrate in the analysis below, it is in this context that our participants are reaching for generative AI tools, using ChatGPT, Claude, and others to address these burdens. It is to these chat data that the article now turns.

Method

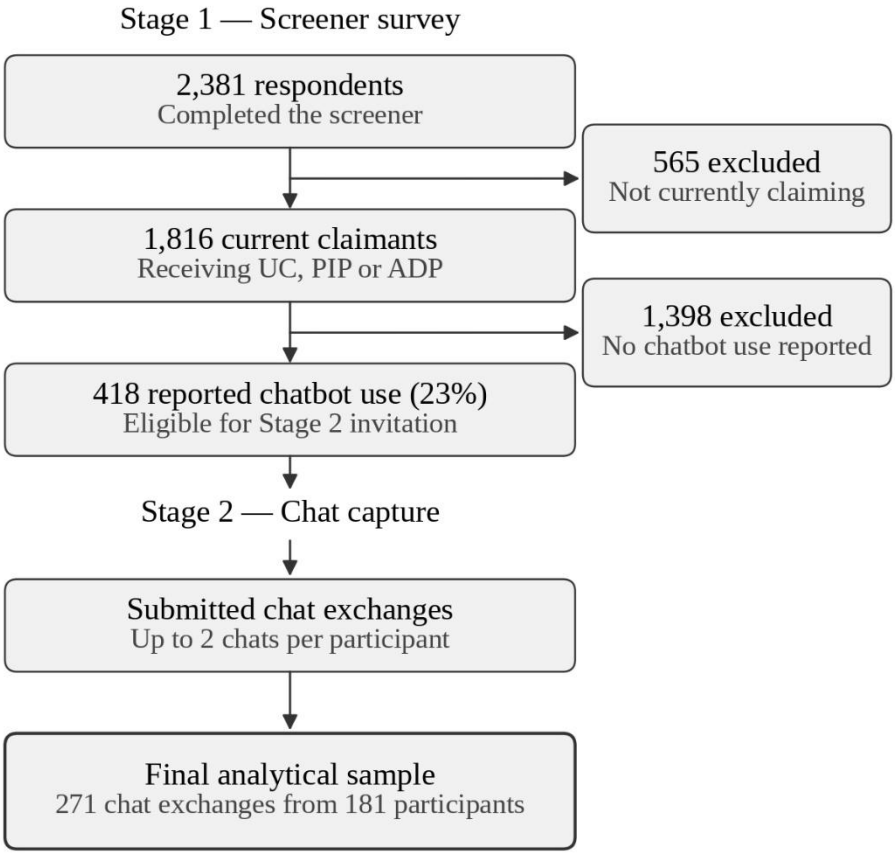
The data analysed in this article were captured through a two-stage method we characterise as 'chat mining': a structured approach to collecting chat prompts that members of the public have entered into publicly available generative AI tools. As outlined in the previous section, most empirical work on generative AI use in public services to date has examined either the outputs of bespoke government-designed chatbots (Cortés-Cediel et al., 2023; Følstad & Bjerkreim-Hanssen, 2024) or hypothetical scenarios of generative AI use (Kieslich et al., 2026). Chat mining goes further by capturing the actual prompts that ordinary users have already submitted to commercial tools such as ChatGPT, Claude, Gemini, Copilot, and Grok while navigating public services.

Both stages of the study were administered through the online research panel Prolific between April and May 2026, a widely used panel provider across the social sciences (for an account of its use elsewhere in this journal, see Harney, 2025, p. 1161). We recruited the sample via Prolific for two reasons. First, it allowed us to draw a sample from members of the panel who had already indicated, through Prolific's screening categories, that they were in receipt of a UK social security payment (either UC or, via prior screening from the research team, PIP, or Adult Disability Payment in Scotland (ADP)). Second, its anonymous infrastructure for participant communication and payment meant that the research team could, at no point, link submitted prompts to any identifying participant details.

Stage one consisted of a short screener survey distributed to members of the Prolific panel identified as currently or previously receiving either UC, PIP/ADP, or both (the survey instruments are provided in the Online Appendix and this process is set out in Figure 1 below). Following a consent process, participants were asked to confirm their current claim status (UC, PIP/ADP, both, or neither) and to indicate which of a set of nine sources they had used for advice or research when interacting with the DWP (ranging from a friend or family member to a solicitor), which included the option ‘Used an AI tool or chatbot (e.g. ChatGPT, Copilot, Gemini, Claude etc)’. Participants were told that the study explored ‘where they turn for help when they need it’ when accessing public services but stopped short of specifying our focus on AI chatbot use to minimise the risk of overreporting. Those who indicated they had used an AI chatbot tool were asked to provide a one-sentence summary of what they had used it for and

whether they would, in principle, be willing to share examples of their prompts in a follow-up study. Of the 2,381 respondents who completed the screener, 1,816 were currently in receipt of UC, PIP/ADP, or both. Within this group, 23 per cent reported having used an AI chatbot to assist with their claim, placing chatbot use just behind contacting Citizens Advice or a similar charity (26.9%), with the most common sources remaining government websites (73.4%) and search engines (51.3%).

Figure 1. Summary of the two-stage recruitment process



We examined whether those who turned to AI chatbots were disproportionately claimants who otherwise went without advice. The screener data suggest the opposite. Among the current recipients, those who reported using an AI chatbot had drawn on more other sources of help, not fewer: they used a mean of 2.82 other advice sources, compared with 2.33 sources among non-users, and only 4.8 per cent of chatbot users had relied on the tool alone. Chatbot users were significantly more likely than non-users to have also used a search engine (66.2% vs 46.6%), an online forum or social media (37.9% vs 28.3%) and were no less likely to have used any source of support, including Citizens Advice or a charity (29.0% vs 26.0%). Chatbot use did not vary significantly by claim type (UC, PIP/ADP, or both).

Stage two invited respondents who had reported AI chatbot use in stage one, and who had indicated a willingness to take part in follow-up research, to share the prompts they had entered into these tools. After a more detailed consent process, participants were asked which chatbot(s) they had used, what they had been hoping the tool would help them with, and were invited to paste in the text of up to two separate chat exchanges, with space provided for an

initial prompt and up to five follow-up prompts per exchange. We focused on the prompts entered by users rather than the replies returned by the chatbot because, as the underlying models change frequently, prompts provide a more durable basis for analysis, and because focusing on prompts minimises the redaction required of respondents, given the length of many chatbot replies. ChatGPT was by some distance the most frequently reported tool (used by 71.7% of stage two participants), followed by Gemini (13.2%), with Claude, Copilot, and Grok each accounting for only a small share of the sample.

The risks associated with handling user-submitted prompts demanded a careful approach to anonymisation. Participants were given clear written guidance before sharing their prompts, including a worked list of the categories of information they were asked to remove or replace with generic placeholders (such as names, addresses, employer details, account numbers, case references and any other potentially identifying detail). Each submitted chat was then independently reviewed by a member of the research team, and any further identifying information was redacted. At this point, we excluded chats which did not relate to a UC or PIP claim - a small number of participants submitted examples concerned with unrelated matters (such as consumer disputes or general administrative queries directed at other public bodies). Where a participant had submitted two chats, each was assessed independently against this criterion. This screening yielded a final analytical sample of 271 chat exchanges drawn from 181 participants. The demographics of this sample are set out in Table 1 below; women are overrepresented in the sample (at 64 per cent), and we do not claim generalisability (either to UC/PIP recipients or even to the subset of those using generative AI chatbots). For researchers seeking to replicate this approach, we would emphasise that chat prompts can contain deeply personal disclosures and, notwithstanding clear instructions not to submit personal identifying information, the manual review of each chat by a researcher proved necessary to remove names and other information in a minority of the sample.

Table 1. Participant demographics for the included chat sample

Age (Years)	
Mean (SD)	34.4 (9.7)
Median	33
Gender	
Female (%)	63.5%
Male (%)	36.5%
Ethnicity	
White	75.7%
Black	9.9%
Mixed	6.6%

Asian	4.4%
Other	2.2%
Prefer not to say	1.2%

Analysis of the final sample of chats proceeded inductively. Members of the research team coded the full sample of prompts to identify themes in what claimants were using the tools for and in how they were framing their requests. These themes were then mapped onto Herd and Moynihan's (2018) tripartite framework of learning, compliance, and psychological costs set out in the previous section, which provides the analytical scaffolding for the findings that follow. Quotations from prompts are reproduced verbatim throughout, with placeholders preserved where the original text had been anonymised by the participant or the research team. The study received ethical approval from [anonymised for review]. Participants were paid on Prolific for their time; the average payment was £2 per submission, based on a completion time of 10 minutes, with additional costs of approximately £0.30 per participant for the screening survey The Prolific panel recruits internationally and permits screening by country, so the approach is not restricted to the UK and is readily adaptable to research on AI chatbot use in other jurisdictions and other areas of public and social service provision.

Analysis

The prompt data reflect both the specific burdens that claimants face during their claim (the forms, messages, applications, complicated entitlement rules and so on) and their efforts to mitigate the associated costs. Across the 271 chat exchanges in our sample, claimants drew on generative AI tools to address each of the learning, compliance, and psychological costs set out above. As Herd and Moynihan (2018) make clear, and as the wider literature on administrative burden has consistently found, each of these costs interacts and compounds one another (Bennett et al., 2026, p. 29; Heinrich, 2016). The bleed between them is clear across our sample: often a single exchange will begin as a request for information, develop into the drafting of a form or message to the DWP, and end with an expression of relief or anxiety. However, so that we can address the implications of our findings for the literature on the costs of burdens, we present the analysis below by taking each cost in turn while drawing attention throughout to examples where they converge.

Learning costs

A recurring use of generative AI tools in our sample was to address learning costs: the work of acquiring the information necessary to establish and maintain entitlement. This is unsurprising given both the complexity of UC and PIP (Bennett et al., 2026; Harris, 2013), and the well-documented learning costs they impose on claimants, particularly in digital-first systems like UC (see Adelmant & Raso, 2025, p. 427). However, rather than seeking general information, the prompts often requested detailed and highly personalised responses, tailored to the claimant's own circumstances through back-and-forth conversation.

For instance, Participant 3 sought to understand how their savings would affect their UC payment and followed an initial question about the taper rate with more granular queries about the timing of their assessment period: '*I reported it as a change today (16th) but dated as the*

15th, plus did my income. It has used my savings figure that I reported on 15th Nov, not the one I reported today.' It is not a question about the rules in the abstract, but about how they bear upon the particular sequence of dates and figures that constitute this claimant's assessment period.

Learning-cost uses also extended to claimants checking whether decisions had been made correctly, or anticipating how a future change might affect entitlement. For instance, Participant 76 (Figure 2) asked ChatGPT to confirm whether a non-dependant deduction should still apply following a change in tenancy arrangements, asking whether the deduction remained applicable now that they and the only other adult in the property were joint tenants.

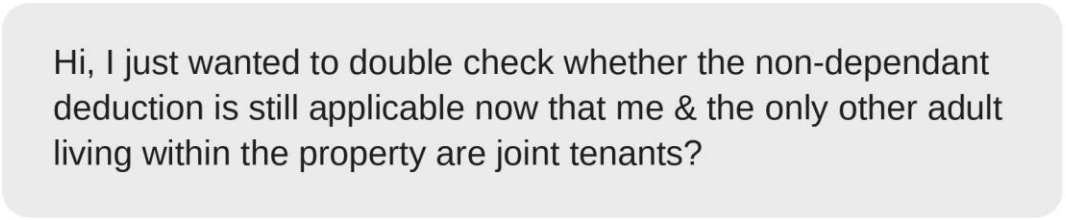


Figure 2. Participant 76

Sometimes, these learning costs were associated with checking the accuracy of a DWP decision. Participant 88 interrogated the calculation of their own award, questioning whether a reduction following a child leaving full-time education had been applied correctly: '*Shouldn't the payment be about £1300 as we have only lost one child's payment which was about £300*'. In these exchanges the tool is used not to explain a rule, but to scrutinise the administration of an individual claim, a form of oversight made difficult by the opacity of automated benefit calculations (see Griffiths, 2024).

Participants also used the tools to navigate the interaction between different sources of support, rather than the rules of a single benefit (see Figure 3). This is the 'multi-level maze' that Bennett (2026) identifies as a defining feature of contemporary social security, in which claimants must hold together entitlements that span benefits, agencies, and tiers of government. Participant 51, for instance, sought to understand how a child becoming eligible for Disability Living Allowance (DLA) would affect the household's UC, before proceeding to questions of reporting, backdating, and the financial implications of different award levels.

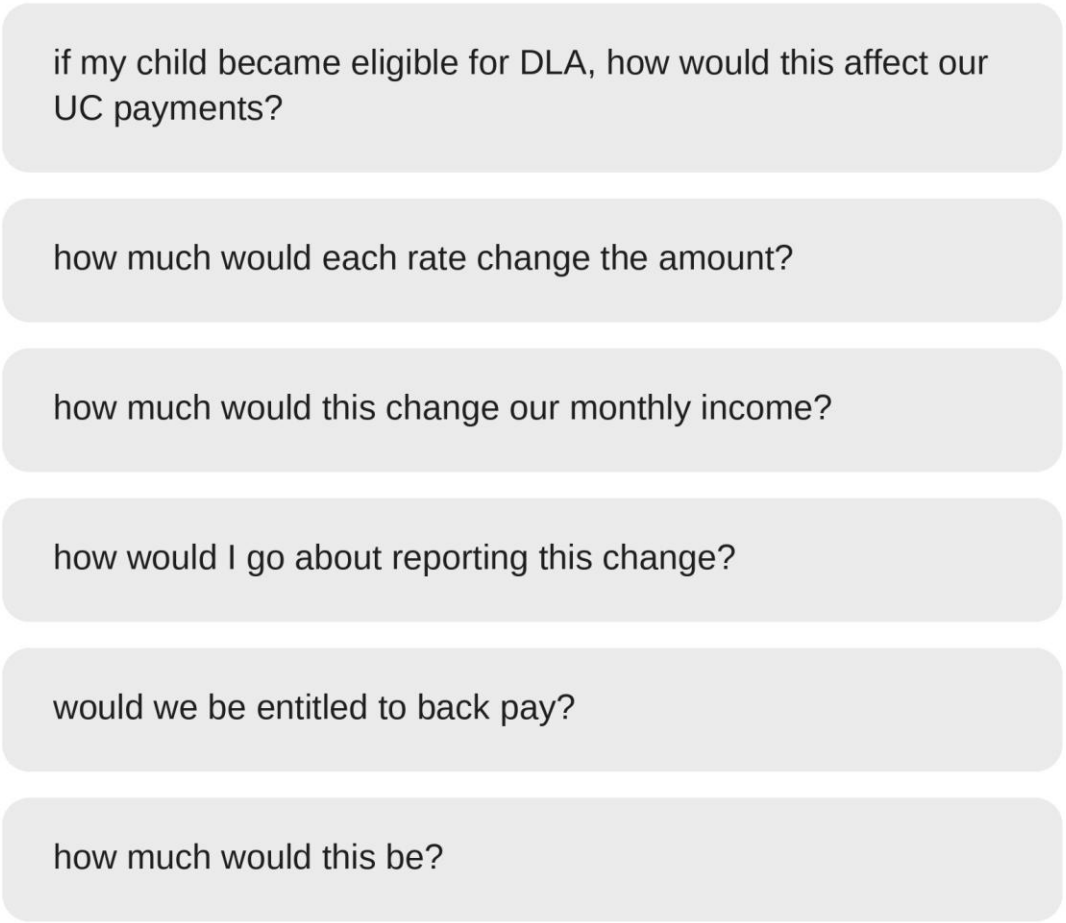


Figure 3. Participant 51

The clearest instances of this took the form of extended exchanges in which a claimant mapped an entire entitlement landscape from the outset. Participant 13's exchange (Figure 4) is especially instructive. Having moved across three distinct benefit schemes within a single thread, they close with a request to rehearse a forthcoming telephone conversation with a doctor, demonstrating how readily the work of understanding the system gives way to the work of managing the anxiety of dealing with it.

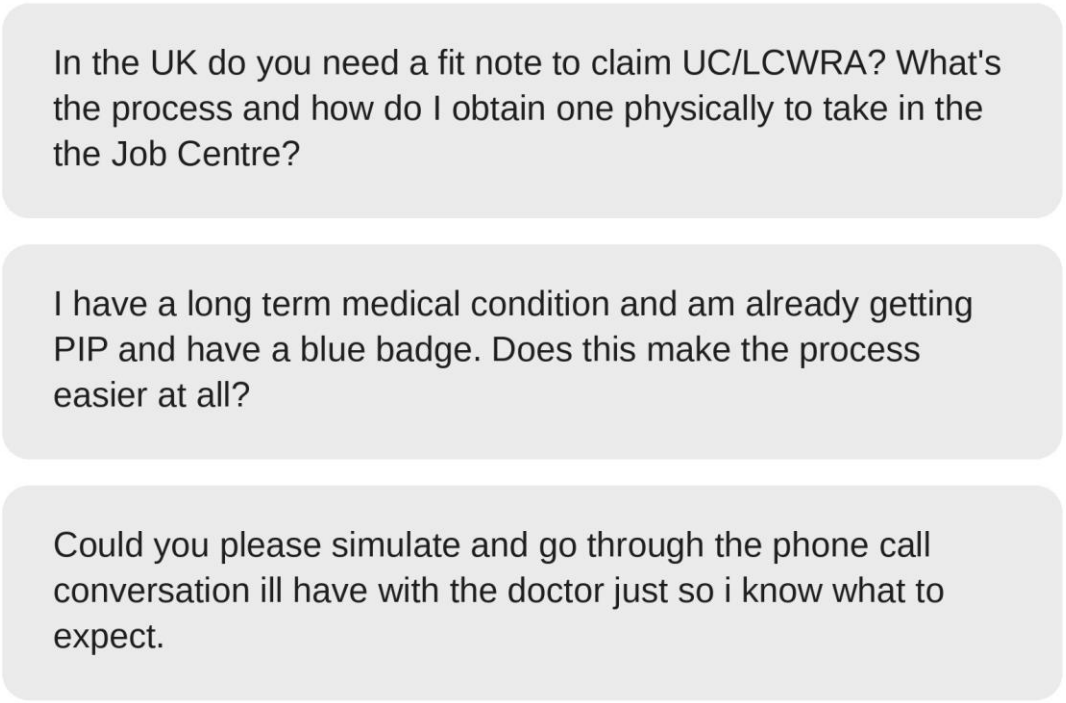


Figure 4. Participant 13

Taken together, the learning-cost exchanges suggest that claimants are using these tools less to discover what the rules are than to establish what the rules mean for them. This highly personalised approach to meeting learning costs is qualitatively different from what is possible via search engines or online guidance, where the burden of translating from a general rule to their particular circumstances falls on the claimant. The tools move instead towards a function closer to that of an adviser. That shift becomes more pronounced still in the compliance cost exchanges to which we now turn.

Compliance costs

If the learning-cost exchanges show claimants establishing what the rules mean, the compliance-cost exchanges show them acting upon that understanding by translating their own circumstances into the language and format that the benefits system demands. Compliance costs are the burdens of demonstrating that entitlement criteria have been met and of fulfilling the obligations attached to receiving support, encompassing form-filling, the provision of evidence, the reporting of changes, etc. (Moynihan et al., 2015). Forms are the archetypal site of these costs, and Brown et al. (2021) have shown how their design can impose substantial burdens (see also Meers, 2024; Ryan, 2023). It is precisely this translation work that claimants in our sample most frequently asked the tools to perform. In doing so, they used the chatbot in ways that resemble the functions traditionally carried out by a welfare adviser, advocate or support worker, or by a well-placed member of a claimant's social network.

The most common form this took was translating a claimant's experience into the terms of a form or descriptor for a PIP application. Participant 8, applying for DLA for a teenage daughter on a waiting list for autism support, described in detail the assistance her daughter required

each day, with dressing, cleaning her teeth, packing a school bag, and getting into school, together with the violent meltdowns that followed if support was not provided, before asking: '*How can I get this across in the forms?*'

For PIP, where entitlement depends on demonstrating the impact of a condition against a set of descriptors (Day & Shaw, 2022), this translation work frequently extended across an entire exchange. Participant 12 (Figure 5) shared detailed descriptions of their physical limitations and asked the tool to tailor their answers to the PIP descriptors, adding further conditions across successive follow-ups before asking what score they might expect.

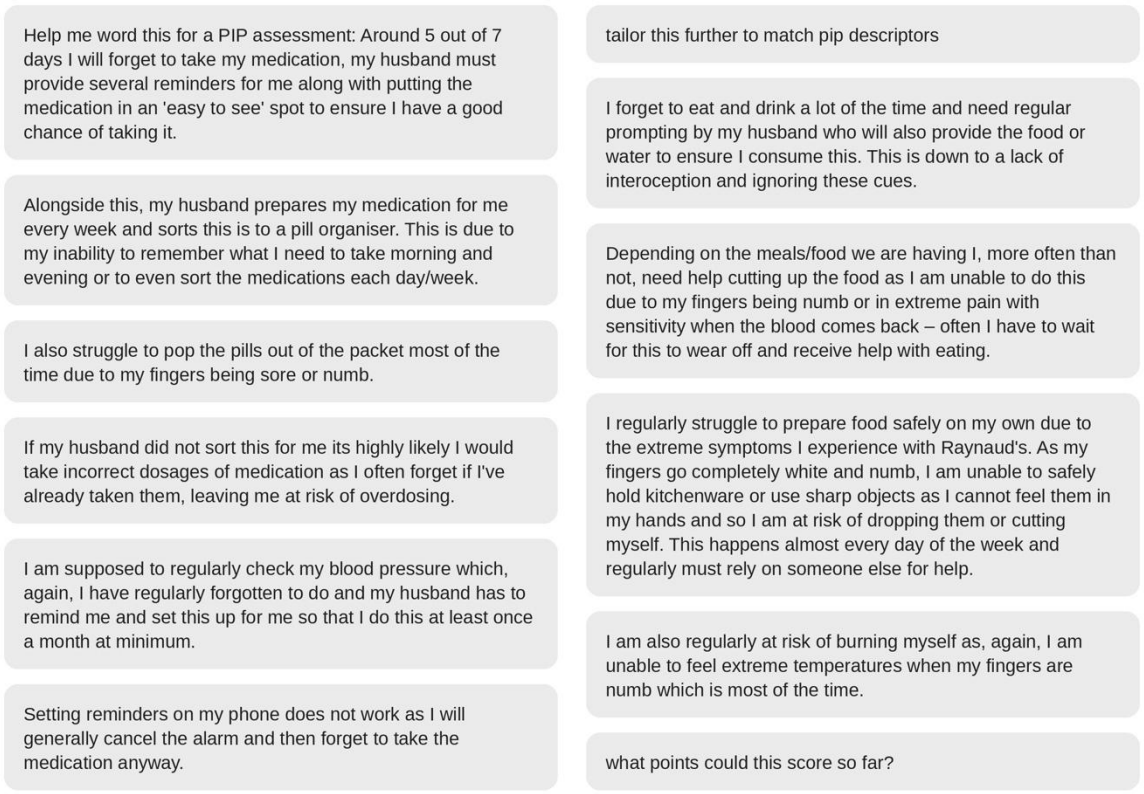


Figure 5. Participant 12

Participant 19 (Figure 6) adopted the same approach on behalf of their son, beginning from a lengthy document describing his autism spectrum disorder (ASD), eating disorder, anxiety, and depression, and progressively refining the drafted answers to capture the fluctuating and overlapping character of those conditions.

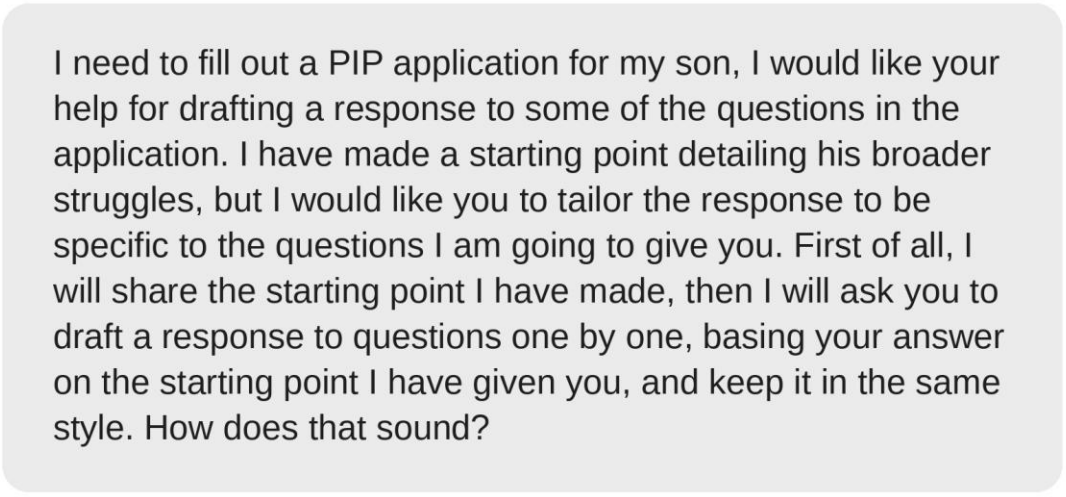


Figure 6. Participant 19

Participant 39 (Figure 7) went further still, using the tool to assess a completed PIP application and asking it to model two readings of the same evidence. In asking the tool to score their application as both a fair and a strict assessor might, and to estimate their prospects at each successive stage, Participant 39 was not seeking to understand the rules but to anticipate how those rules would be applied by a particular decision-maker. Here the claimant is modelling the discretion of the official using the chatbot and adjusting their own submission considering that model.

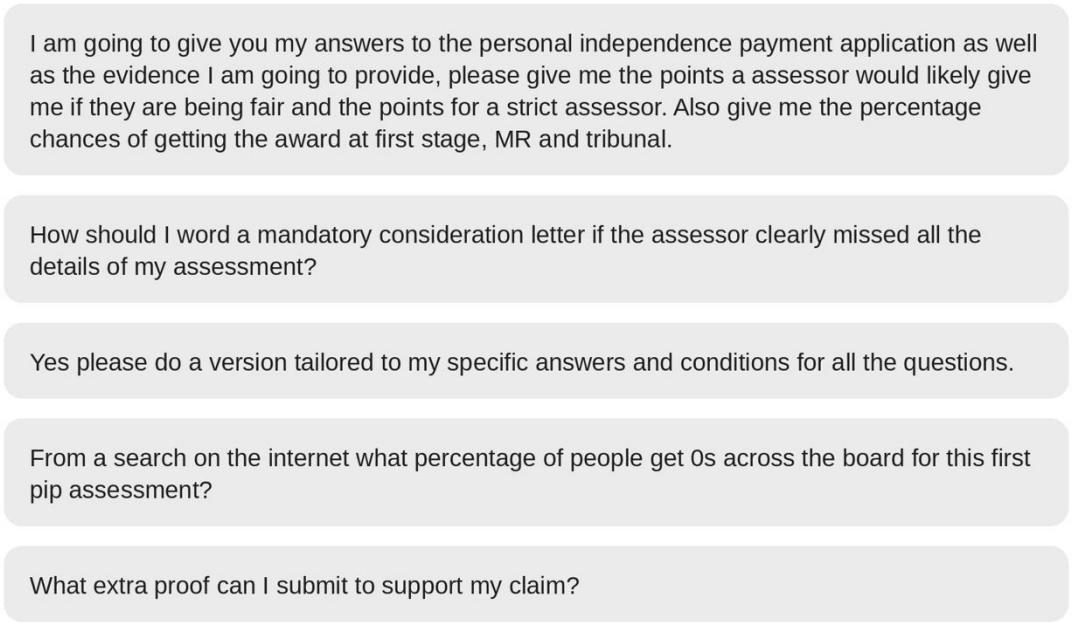


Figure 7. Participant 39

Beyond the form itself, claimants used the tools to manage the continuous compliance obligations that characterise UC, where burdens do not conclude at application but persist

across the life of the claim (Meers et al., 2025). For instance, Participant 43 (Figure 8) asked the tool to draft a journal message to the DWP explaining that Employment and Support Allowance (ESA) had not been paid owing to a late sick note, so that UC would not deduct it pound for pound.

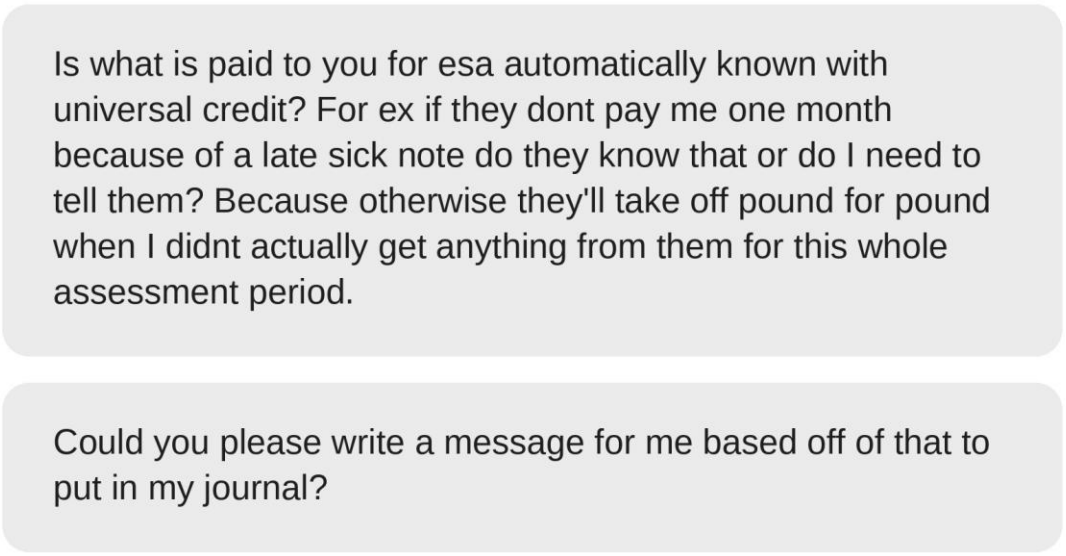


Figure 8. Participant 43

Participant 57 used the tool both to understand how to apply for an advance payment and to compose the message requesting it, asking: '*How do I go about asking for this? Can you draft me a journal message for my work coach?*'. In each case, these participants are using the chatbot to draft required text, whether the journal entry or the message to a work coach, through which compliance with the system of benefit rules is achieved.

The most striking compliance-cost uses, however, concerned not making a claim but challenging a decision. Participant 95 (Figure 9) asked the tool to '*Write me a reconsideration appeal for a UC sanction*', before requesting that the draft be made shorter, less formal, and more persuasive.

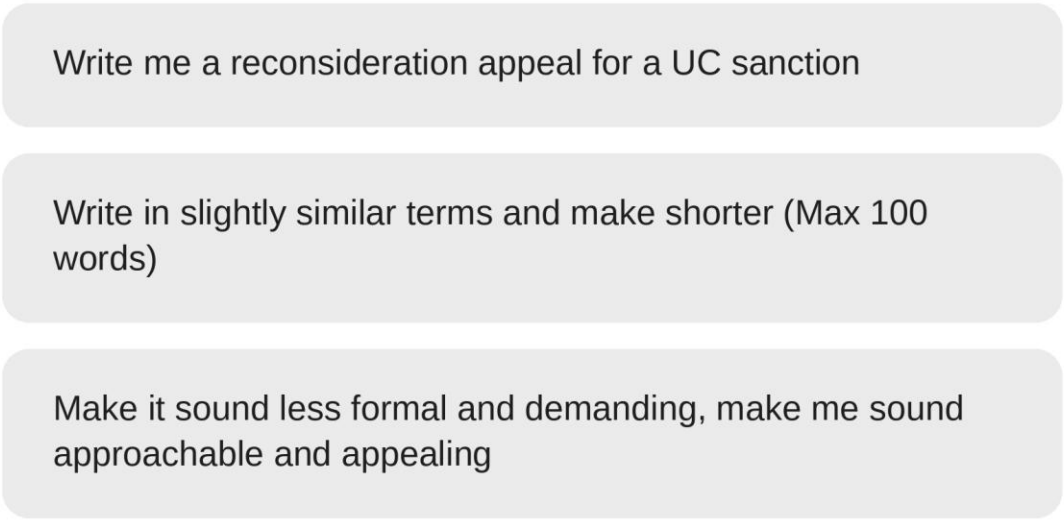


Figure 9. Participant 95

As these excerpts illustrate, claimants in the sample were using these tools to draft correspondence of all kinds, including complex interactions requiring significant additional context, such as the internal review of decisions. The use of generative AI to draft challenges to benefit decisions may generate additional demand for redress processes, an issue we return to in the discussion and conclusion.

Psychological costs

Psychological costs, namely the stress, anxiety, frustration, stigma and loss of autonomy that can accompany dealings with the welfare state, run throughout the chat data and are the most difficult of the three to isolate (Bennett et al., 2026, p. 37). They are the least visible of the burdens, yet some literature suggests they may be the most consequential, shaping both whether people claim at all and how they experience claiming (Baumberg, 2016; Patrick, 2017). The well-documented psychological toll of UC and PIP, comprising the stress of waiting for a decision, the anxiety of assessment and reassessment, the fear of sanctions, and the stigma of claiming, surfaces repeatedly in our data (Cheetham et al., 2019; Dwyer et al., 2020). What the chat exchanges add is a view of how claimants seek to use these tools to manage these burdens.

In some exchanges the psychological cost was the explicit subject of the interaction. Participant 44, applying for PIP, set out what they intended to include and then, midway through, disclosed their anxiety in the chat directly: '*I am really stressed and anxious I won't be able to finish and get accepted. Is it rare for people like me to get accepted?*' Participant 29, who described using the tool '*usually for reassurance more than anything*', rehearsed and validated their position ahead of a work coach meeting before asking: '*And if it all goes terribly wrong, will you be here to talk through it with me afterwards?*' In both cases the tool functioned less as a source of information than as a source of reassurance in anticipation of an encounter with the system.

More striking still are the exchanges in which psychological costs determined whether the learning and compliance work was undertaken at all. This is the compounding dynamic that the burdens literature describes (Bennett et al., 2026; Heinrich, 2016), captured in real time.

Participant 40 described a '*mental block*' that had lasted several months, explaining that '*now there's sanctions and stuff on my universal credit. I don't know what to write in my journal*'. The obstacle here was not a lack of information but the burden of re-engaging with a system after falling behind. Participant 42 (Figure 10) sought help after realising they had failed to declare self-employment income, stating that they were '*shitting myself at the thought that I have done something wrong*'. That exchange then developed into a detailed discussion of reporting requirements and record-keeping before the claimant returned to report the outcome.

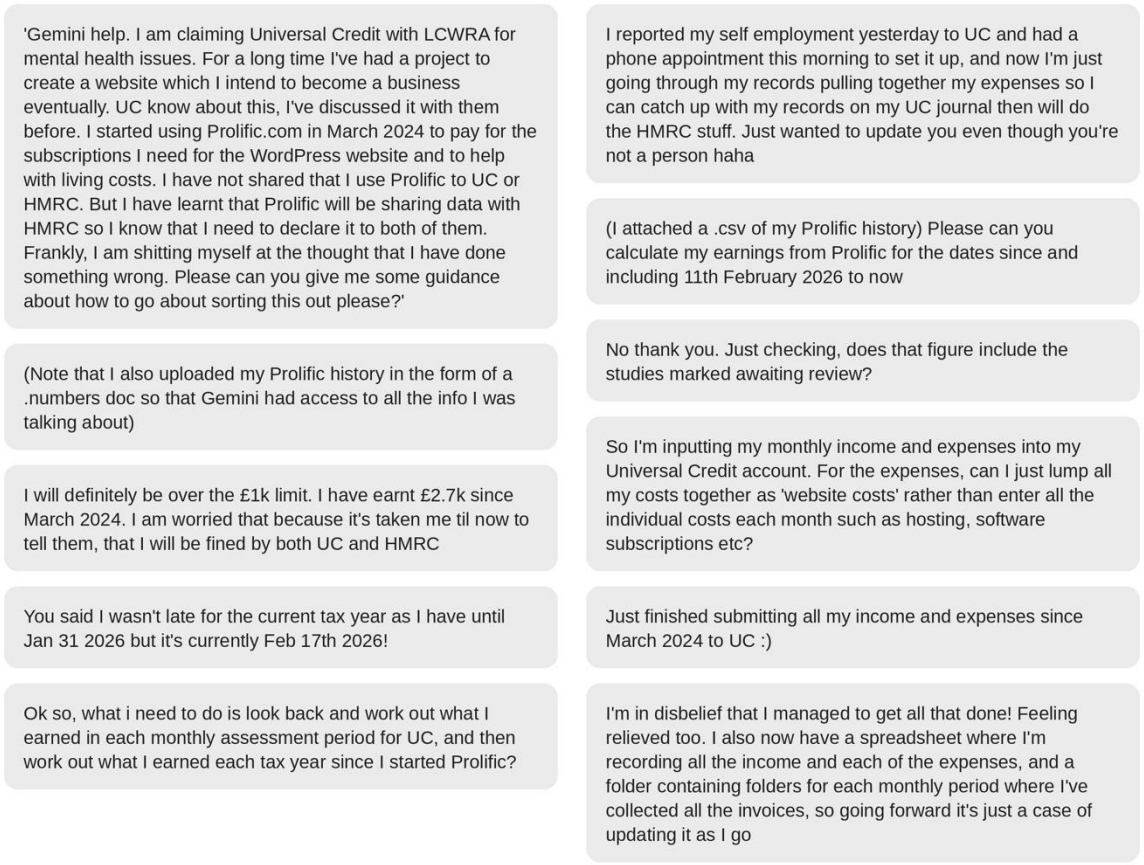


Figure 10. Participant 42

In other exchanges the psychological dimension surfaced more indirectly, in the tone of the prompts themselves. Participant 9, bedbound and describing themselves as struggling to write, sought help with a PIP application before disclosing both a fear of being misunderstood and, revealingly, that they '*don't want to get in trouble for using ChatGPT*'. Participant 51, having asked whether a work coach could demand confirmation emails of job applications, pushed back against the tool itself: '*Is this absolutely true? I want facts and not you being sycophantic.*' Participant 98, reflecting on disability assessments, observed that '*It's designed to make people give up.*' The psychological costs of claiming, then, frequently determine whether that other work is attempted at all. The use of these AI chatbots appears to have lowered the barrier by offering an accessible, responsive, and seemingly non-judgemental point of contact, available at the moment the burden was felt and without the stigma that claimants may associate with other sources of help. This echoes experimental evidence that people disclose more readily to machine interlocutors, precisely because they feel less judged by them (Lucas et al., 2014).

These findings echo some of the emerging literature on conversational AI as a source of companionship and emotional support (for an influential contribution to this burgeoning literature, see Skjuve et al., 2021). When Participant 29 asks whether the tool will be there to talk things through after a difficult meeting, its use is not only to offset a burden imposed by the state at a discreet point in time, but also to provide an ongoing relational resource for coping with the experience of being a claimant.

Discussion and Conclusion

Our dataset of chat prompts shows that claimants are using widely available generative AI tools (namely ChatGPT) to help meet the costs of administrative burdens that come with making and maintaining a claim for benefits in the UK. Given the increasing prevalence and adoption of these tools across society (Department for Science, Innovation and Technology, 2026), it is perhaps unsurprising to see their use bleed into interactions with the social security system. We suggest that the findings in this article, and more broadly the chat prompts from claimants as a dataset for analysis, are significant for at least four reasons.

First, as flagged in our discussion of the method underpinning this paper, chat prompt data provide a valuable and hitherto underutilised insight into the burdens faced by people when they seek to access public services. In the wide-ranging field of research on administrative burden, the methodological focus tends to be either on data collected at the point of interaction with the state (for example, measuring the experience of applying for a benefit via a survey) or on post-hoc accounts. Chat prompt data provide an immediate snapshot of what the individual was struggling with at a particular point in time, written in their own words. Other work, particularly in economics, has drawn on search engine query data (such as Google Trends; for influential work in this space, see Choi & Varian, 2012), but our data suggest that chatbot prompts provide a far richer and more revealing source: individuals often have extended back-and-forth conversations, share documents and reveal their anxieties and concerns. Our approach to collecting chat prompts from UC and PIP recipients could therefore be extended across other areas of the state: health and social care, local government, immigration, and so on.

Second, our results suggest that claimants are using generative AI tools to support them in a wide range of interactions with the state, from writing journal messages in the UC interface, to drafting application forms for benefits, to lodging appeals or complaints. The prompts in our dataset are, in effect, claimants reaching for AI tools at the points where the system's learning, compliance, and psychological costs press hardest, and using those tools to try to offset them. What this means for public services as the tools proliferate is, we think, difficult to predict. The most obvious consequence is more volume: more text, more applications, more appeals that might not otherwise have been pursued. But there are other possibilities too. AI tools could increasingly decouple the text from its author, with claimants making applications or sending messages not just that they may not have written, but also that they may not have even read (for a civil servant-facing view of this issue, see Meers, 2026). Errors may find their way into decision-making in new ways, and operational assumptions that the front-line of public administration has long taken for granted may need rethinking. As agent functionality develops further, this may generate even more disruption, cause even more demand, and disrupt administrative systems long predicated on human-inputted text. We do not yet know, and that

is part of the point: this is an unsettled moment, and public administration will have to work out the implications for the design and delivery of services.

Third, the findings extend the conceptual architecture of the administrative burden framework itself. Claimants have never borne burdens entirely alone: friends and family, welfare rights advisers, and organisations such as Citizens Advice have long acted as intermediaries who absorb part of the work of learning and compliance (Friskney et al., 2026). Our data show a new intermediary entering that landscape of support, but these are commercial products, developed by a small number of technology firms, freely available at any hour, and changing in capability with each model release. The costs of burdens they help to meet are consequently mediated by tools whose design reflects the priorities of technology firms rather than of welfare agencies or the advice sector. In this way, the question the framework poses, of who bears the costs of administration, now turns in part on the kind of intermediary through which those costs are borne. Where claimants use these tools for reassurance, rehearsal, and emotional support, the chatbot is not so much reducing a cost imposed by the state as supplying a resource for living with it, a use closer to the developing literature on AI companionship (Skjuve et al., 2021) than to established accounts of burden mitigation. Nor is this new intermediary available to everyone. Many claimants remain digitally excluded, lacking the access, skills or confidence to use these tools at all, so the capacity to offset burdens in this way may itself be unevenly distributed and concentrated among those already best placed to seek help (Alexopoulou, 2025; Wang et al., 2025).

Fourth, unlike assistant chatbots that public services design and develop themselves (see Følstad & Bjerkreim-Hanssen, 2024; Verne et al., 2022), these tools sit beyond the sight and control of public services. ChatGPT, Claude, and Gemini are free to access and regularly updated. Much of the current literature and wider debate asks how the welfare state does, could, or should use AI. Our findings point to a different dynamic: AI is already changing the welfare state from the other direction, through how claimants are using it, and social policy and administration will have to respond. Disruption, in this view, does not come from the state alone; it can come from the public side of the interface as well.

There is, in turn, a substantial agenda for social policy and administration research. We do not yet know how accurate the guidance these tools return to claimants is, nor how officials, advisers, and others will react as the scale of claimant use becomes apparent. Nor do we know how these dynamics will play out beyond the UK, in welfare systems with different administrative cultures and degrees of digitalisation, both in the state and across the population. What is clear is that the pressures facing welfare administrators, and the character of citizen-state interactions, are being reshaped by claimants themselves at a pace that scholarship has yet to register. Understanding the scale and consequences of that change will require sustained empirical attention, and the method set out in this article offers one way of providing it.

References

Adelmant, V., & Raso, J. (2025). Data Entry and Decision Chains: Distributed Responsibility and Bureaucratic Disempowerment in the UK's Universal Credit Programme. *Oxford Journal of Legal Studies*, *45*(2), 415–446. https://doi.org/10.1093/ojls/gqaf006

Agostino, D., Saliterer, I., & Steccolini, I. (2022). Digitalization, accounting and accountability: A literature review and reflections on future research in public

services. *Financial Accountability & Management*, *38*(2), 152–176. https://doi.org/10.1111/faam.12301

Alexopoulou, S. (2025). Paradigm shift: Exploring the impact of digital technologies on the welfare state through a systematic literature review. *Social Policy & Administration*, *59*(1), 135–157. https://doi.org/10.1111/spol.13054

Alishani, A., & Homburg, V. (2025). When citizens meet the chatbot: Evidence from a survey vignette experiment in Estonia. *Public Policy and Administration*, 09520767251404286. https://doi.org/10.1177/09520767251404286

Androutsopoulou, A., Karacapilidis, N., Loukis, E., & Charalabidis, Y. (2019). Transforming the communication between citizens and government through AI-guided chatbots. *Government Information Quarterly*, *36*(2), 358–367. https://doi.org/10.1016/j.giq.2018.10.001

Aoki, N. (2020). An experimental study of public trust in AI chatbots in the public sector. *Government Information Quarterly*, *37*(4), 101490. https://doi.org/10.1016/j.giq.2020.101490

Baumberg, B. (2016). The stigma of claiming benefits: A quantitative study. *Journal of Social Policy*, *45*(2), 181–199. https://doi.org/10.1017/S0047279415000525

Bennett, H. (2026). Navigating Digital Welfare: A Multi-Level Maze? *Social Policy & Administration*, *60*(1), 27–38. https://doi.org/10.1111/spol.13138

Bennett, H., Currie, M., & Podoletz, L. (2026). Universal Credit: Administrative burdens of automated welfare. *Journal of Social Policy*, *55*(1), 24–42. https://doi.org/10.1017/S0047279424000175

Bright, J., Enock, F., Esnaashari, S., Francis, J., Hashem, Y., & Morgan, D. (2025). Generative AI is already widespread in the public sector: Evidence from a survey of UK public sector professionals. *Digital Government: Research and Practice*, *6*(1), 2:1-2:13. https://doi.org/10.1145/3700140

Brown, J. T., Carey, G., & Malbon, E. (2021). What is in a form? Examining the complexity of application forms and administrative burden. *Australian Journal of Public Administration*, *80*(4), 933–964. https://doi.org/10.1111/1467-8500.12531

Capraro, V., Lentsch, A., Acemoglu, D., Akgun, S., Akhmedova, A., Bilancini, E., Bonnefon, J.-F., Brañas-Garza, P., Butera, L., Douglas, K. M., Everett, J. A. C., Gigerenzer, G., Greenhow, C., Hashimoto, D. A., Holt-Lunstad, J., Jetten, J., Johnson, S., Kunz, W. H., Longoni, C., … Viale, R. (2024). The impact of generative artificial intelligence on socioeconomic inequalities and policy making. *PNAS Nexus*, *3*(6), pgae191. https://doi.org/10.1093/pnasnexus/pgae191

Cheetham, M., Moffatt, S., Addison, M., & Wiseman, A. (2019). Impact of Universal Credit in North East England: A qualitative study of claimants and support staff. *BMJ Open*, *9*(7), e029611. https://doi.org/10.1136/bmjopen-2019-029611

Chien, C. V., & Kim, M. (2024). *Generative AI and Legal Aid: Results from a Field Study and 100 Use Cases to Bridge the Access to Justice Gap* (SSRN Scholarly Paper No. 4733061). Social Science Research Network. https://papers.ssrn.com/abstract=4733061

Choi, H., & Varian, H. (2012). Predicting the Present with Google Trends. *Economic Record*, *88*(s1), 2–9. https://doi.org/10.1111/j.1475-4932.2012.00809.x

Choi, S., Sung, W., & Lee, J. (2025). Exploring generative AI effects on Korean civil servants' performance: Power and motivation. *Conference on Digital Government Research*, *26*. Conference on Digital Government Research. https://doi.org/10.59490/dgo.2025.1032

Considine, M., Mcgann, M., Ball, S., & Nguyen, P. (2022). Can Robots Understand Welfare? Exploring Machine Bureaucracies in Welfare-to-Work. *Journal of Social Policy*, *51*(3), 519–534. https://doi.org/10.1017/S0047279422000174

Cortés-Cediel, M. E., Segura-Tinoco, A., Cantador, I., & Rodríguez Bolívar, M. P. (2023). Trends and challenges of e-government chatbots: Advances in exploring open government data and citizen participation content. *Government Information Quarterly*, *40*(4), 101877. https://doi.org/10.1016/j.giq.2023.101877

Day, W., & Shaw, R. (2022). When benefit eligibility and patient-led care intersect. Living in the UK with chronic illness: Experiences of the work capability assessment. *Journal of Health Psychology*, *27*(2), 456–469. https://doi.org/10.1177/1359105320953476

Department for Science, Innovation and Technology. (2026). *AI Skills for Life and Work: General Public Survey Findings*. GOV.UK. https://www.gov.uk/government/publications/ai-skills-for-life-and-work-general-public-survey-findings/ai-skills-for-life-and-work-general-public-survey-findings

Department for Work and Pensions. (2026a). *Artificial Intelligence Security Policy*. GOV.UK. https://www.gov.uk/government/publications/dwp-procurement-security-policies-and-standards/artificial-intelligence-security-policy

Department for Work and Pensions. (2026b). *DWP benefit statistics: February 2026*. GOV.UK. https://www.gov.uk/government/statistics/dwp-benefit-statistics-february-2026

Department for Work and Pensions. (2026c). *An evaluation of DWP's Microsoft 365 Copilot trial*. GOV.UK. https://www.gov.uk/government/publications/an-evaluation-of-dwps-microsoft-copilot-365-trial/an-evaluation-of-dwps-microsoft-365-copilot-trial

Dwyer, P., Scullion, L., Jones, K., McNeill, J., & Stewart, A. B. R. (2020). Work, welfare, and wellbeing: The impacts of welfare conditionality on people with mental health impairments in the UK. *Social Policy & Administration*, *54*(2), 311–326. https://doi.org/10.1111/spol.12560

Følstad, A., & Bjerkreim-Hanssen, N. (2024). User Interactions With a Municipality Chatbot—Lessons Learnt From Dialogue Analysis. *International Journal of Human–Computer Interaction*, *40*(18), 4973–4986. https://doi.org/10.1080/10447318.2023.2238355

Friskney, R., Halliday, S., Macintyre, G., Meers, J., Seyd, B., & Tomlinson, J. (2026). The Importance of Assistance for Disability Benefits Claiming: Administrative Burden, Bureaucratic Rationality, and the False Promise of Equality. *Journal of Social Security Law*, *33*(2).

Griffiths, R. (2024). Universal Credit and Automated Decision Making: A Case of the Digital Tail Wagging the Policy Dog? *Social Policy and Society*, *23*(1), 1–18. https://doi.org/10.1017/S1474746421000749

Harney, J. (2025). Expanding Perceived Capacity of Public Sector Workers: Experimental Evidence on the Impact of a Counter-Stereotypical Narrative on Public Perceptions of Correctional Officers. *Social Policy & Administration*, *59*(7), 1158–1168. https://doi.org/10.1111/spol.13118

Harris, N. (2013). *Law in a Complex State*. Hart Publishing. https://www.bloomsbury.com/uk/law-in-a-complex-state-9781849464451/

Heinrich, C. J. (2016). The Bite of Administrative Burden: A Theoretical and Empirical Investigation. *Journal of Public Administration Research and Theory*, *26*(3), 403–420. https://doi.org/10.1093/jopart/muv034

Herd, P., & Moynihan, D. P. (2018). *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation. https://doi.org/10.7758/9781610448789

Jørgensen, R. F. (2023). Data and rights in the digital welfare state: The case of Denmark. *Information, Communication & Society*, *26*(1), 123–138. https://doi.org/10.1080/1369118X.2021.1934069

Kaun, A., & Männiste, M. (2025). Public sector chatbots: AI frictions and data infrastructures at the interface of the digital welfare state. *New Media & Society*, *27*(4), 1962–1985. https://doi.org/10.1177/14614448251314394

Kieslich, K., Helberger, N., & Diakopoulos, N. (2026). *"Make It Sound Like a Lawyer Wrote It": Scenarios of Potential Impacts of Generative AI for Legal Conflict Resolution* (arXiv:2602.24130). arXiv. https://doi.org/10.48550/arXiv.2602.24130

Larsen, A. G., & Følstad, A. (2024). The impact of chatbots on public service provision: A qualitative interview study with citizens and public service providers. *Government Information Quarterly*, *41*(2), 101927. https://doi.org/10.1016/j.giq.2024.101927

Lucas, G. M., Gratch, J., King, A., & Morency, L.-P. (2014). It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, *37*, 94–100. https://doi.org/10.1016/j.chb.2014.04.043

Madsen, C. Ø., Lindgren, I., & Melin, U. (2022). The accidental caseworker – How digital self-service influences citizens' administrative burden. *Government Information Quarterly*, *39*(1), 101653. https://doi.org/10.1016/j.giq.2021.101653

Meers, J. (2024). Discrimination at the Interface: The Equality Act 2010 and Platform Interface Design. *The Modern Law Review*, *87*(3), 640–669. https://doi.org/10.1111/1468-2230.12855

Meers, J. (2026). 'Chucking Their Stuff into ChatGPT and Hitting Send': AI Slop in the Administrative State. *Journal of Law and Society*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6898138

Meers, J., Halliday, S., & Tomlinson, J. (2025). An 'interface first' bureaucracy: Interface design, universal credit and the digital welfare state. *Social Policy & Administration*, *59*(1), 119–134. https://doi.org/10.1111/spol.13053

Moynihan, D., Herd, P., & Harvey, H. (2015). Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions. *Journal of Public Administration Research and Theory*, *25*(1), 43–69. https://doi.org/10.1093/jopart/muu009

OpenAI. (2026, May 20). *How people are using ChatGPT*. OpenAI. https://openai.com/index/how-people-are-using-chatgpt/

Patrick, R. (2017). *For Whose Benefit?* Policy Press.

Peeters, R. (2020). The Political Economy of Administrative Burdens: A Theoretical Framework for Analyzing the Organizational Origins of Administrative Burdens. *Administration & Society*, *52*(4), 566–592. https://doi.org/10.1177/0095399719854367

Pors, A. S., & Pallesen, E. (2021). The reorganization of the bureaucratic encounter in a digitized public administration. *Ephemera: Theory & Politics in Organization*, *21*(3), 17–41.

Rachovitsa, A., & Johann, N. (2022). The Human Rights Implications of the Use of AI in the Digital Welfare State: Lessons Learned from the Dutch SyRI Case. *Human Rights Law Review*, *22*(2), ngac010. https://doi.org/10.1093/hrlr/ngac010

Ryan, A. (2023). The Form of Forms: Everyday Enablers of Access to Justice. *Social & Legal Studies*, 32(5), 690–713. https://doi.org/10.1177/09646639231172616

Ryan, F., & Hardie, L. (2024). ChatGPT, I have a Legal Question? The Impact of Generative AI Tools on Law Clinics and Access to Justice. *International Journal of Clinical Legal Education*, *31*(1), 166–205. https://doi.org/10.19164/ijcle.v31i1.1401

Schou, J., & Pors, A. S. (2019). Digital by default? A qualitative study of exclusion in digitalised welfare. *Social Policy & Administration*, *53*(3), 464–477. https://doi.org/10.1111/spol.12470

Senadheera, S., Yigitcanlar, T., Desouza, K. C., Mossberger, K., Corchado, J., Mehmood, R., Li, R. Y. M., & Cheong, P. H. (2025). Understanding Chatbot Adoption in Local Governments: A Review and Framework. *Journal of Urban Technology*, *32*(3), 35–69. https://doi.org/10.1080/10630732.2023.2297665

Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2021). My chatbot companion – A study of human-chatbot relationships. *International Journal of Human-Computer Studies*, *149*, 102601. https://doi.org/10.1016/j.ijhcs.2021.102601

Somers-Joce, C., & Tomlinson, J. (2026). Human rights, public law, and administrative burden: In the matter of an application by JR87 and another for judicial review. *The Modern Law Review*. Advance online publication. https://doi.org/10.1111/1468-2230.70032

UK Government. (2024). *Guidance to civil servants on use of generative AI*. GOV.UK. https://www.gov.uk/government/publications/guidance-to-civil-servants-on-use-of-generative-ai/guidance-to-civil-servants-on-use-of-generative-ai

United Nations. (2022). *United Nations E-Government Survey 2022: The Future of Digital Government*. United Nations Department of Economic and Social Affairs. https://publicadministration.un.org/egovkb/en-us/reports/un-e-government-survey-2022

United Nations. (2024). *United Nations E-Government Survey 2024: Accelerating Digital Transformation for Sustainable Development*. United Nations Department of Economic and Social Affairs. https://publicadministration.un.org/egovkb/en-us/Reports/UN-E-Government-Survey-2024

van Noordt, C., & Misuraca, G. (2019). *New Wine in Old Bottles: Chatbots in Government: Exploring the Transformative Impact of Chatbots in Public Service Delivery* In *Electronic Participation* (pp. 49–59). Springer. https://doi.org/10.1007/978-3-030-27397-2_5

Verne, G., Steinstø, T., Simonsen, L., & Bratteteig, T. (2022). How Can I Help You? A chatbot's answers to citizens' information needs. *Scandinavian Journal of Information Systems*, *34*(2). https://aisel.aisnet.org/sjis/vol34/iss2/7

Wang, C., Boerman, S. C., Kroon, A. C., Möller, J., & de Vreese, C. H. (2025). The artificial intelligence divide: Who is the most vulnerable? *New Media & Society*, *27*(7), 3867–3889. https://doi.org/10.1177/14614448241232345