

Accommodation Repricing:

Disability, Generative AI Subscriptions, and the Accessibility Tax as a Variable the Vendor

Controls

Travis Gilly 🦋 August 2026

ORCID: 0009-0007-2954-6313

Real Safety AI Foundation

Illinois, United States

ABSTRACT

Disabled people increasingly meet access needs with general purpose generative artificial intelligence sold as a commercial subscription rather than with prescribed assistive technology. Three findings are established in the accessible computing literature: the use is widespread across disability communities, it is frequently necessary rather than discretionary, and pricing tiers determine the quality of access a disabled user can obtain. This paper extends the third finding from level to change. Prior work describes what a disabled user can afford at a moment in time. This paper describes what happens when the price, the metered quantity, or the capability itself moves after the user has already reorganized around it, and it names that event accommodation repricing. The argument is that two literatures each hold half of the mechanism and neither holds the join. The accessibility tax literature describes the standing cost of inaccessibility and treats that cost as a level. The assistive technology abandonment literature explains why access tools fall out of use but locates the cause in the device, in the match between device and user, or in the user's own changing needs; a complete reading of the two most cited taxonomies in that literature finds no category at all for a cost the supplier moves after adoption. The paper identifies three positions a disabled user can occupy with respect to an access capability, and argues the third is the most exposed: prescribed assistive technology, where obligations attach; a consumer subscription, where the user is a weak party to a contract they cannot negotiate; and a platform procured by an institution, where the capability can be metered or withdrawn in a transaction the disabled user is not a party to at all. The paper distinguishes accommodation repricing from technoableism, which explains how the accommodation came to be a commercial product but does no work on what happens afterward, and it takes up the strongest existing objection, that recent power aware work in accessible computing has already displaced the adoption-as-fit account. It declines to coin a new category of ableism on the ground that the mechanism is already named at both ends by existing literatures. It then builds and runs the instrument capable of measuring the phenomenon without asking disabled people to narrate a cost increase they cannot escape: a longitudinal documentary audit of published vendor pricing artifacts, coded from dated public archive captures. A first pass across 16 vendor surfaces, with one consumer surface coded in full across 12 consecutive quarters, finds the predicted behavior in a form stronger than a price increase. The posted price of the paid tier did not change across three years. The quantity behind it was metered, then re-expressed in a unit not commensurable with the one it replaced, leaving the user without the comparison required to detect that anything had been reduced.

Keywords: accommodation repricing; accessibility tax; disability tax; generative artificial intelligence; assistive technology abandonment; subscription pricing; switching costs; technoableism; cript technoscience; educational technology procurement; accessible computing; documentary audit

I. INTRODUCTION

A disabled person who uses a commercial generative artificial intelligence subscription to draft correspondence, transcribe speech, summarize an inaccessible document, or rehearse a difficult conversation is doing

something the accessible computing literature has now documented at length. What that literature has not described is what happens to that person when the subscription changes.

The change arrives in several forms. The monthly price rises. The capability moves from the tier the user is on to a tier above it. A flat allowance becomes a metered credit. The metered credit shrinks against an allowance the vendor has never published. The capability is withdrawn from the consumer product entirely and offered only through an enterprise agreement. In each case the user's access need is unchanged, the user's own conduct is unchanged, and the thing that has moved is a term set by a party the user cannot negotiate with.

This paper names that event accommodation repricing and argues that it falls into a gap between two literatures that each hold half of the mechanism.

The first is the accessibility tax literature, which names the direct and indirect financial costs disabled people pay because the surrounding environment is inaccessible (Olsen, Cork, and Anders 2022). That literature has recently been extended to generative artificial intelligence, where researchers found that current pricing models create tiered accessibility in which financial resources determine the quality of access support available to disabled users (Johnson et al. 2026). That finding is a level. It describes what a person can obtain at a given price at a given moment.

The second is the assistive technology abandonment literature, which asks why access tools fall out of use. It is a large and durable literature with a clear origin, and its explanatory categories have been consistent across three decades: the device performed poorly, the device was a poor match for the person, the person was not consulted in the selection, the person's needs changed, the training or support was absent. Every one of those categories locates the cause either in the artifact or in the user. None of them describes a supplier who changes the terms after the user has adapted.

Put the two together and the gap is visible. The accessibility tax literature has the cost but treats it as static. The abandonment literature has the loss of the tool but attributes the loss to the tool or to the person. Accommodation repricing is a cost that moves, imposed by a party neither literature models, producing a loss of access that neither literature can classify.

The contribution is that single move, from level to change. Everything else follows from it. The paper states the pricing mechanism briefly, because it is well described in economics and is not this paper's subject. It then identifies three positions a disabled user can occupy with respect to an access capability and shows that the protections attaching to the first do not travel into the second or the third. It shows why the abandonment taxonomies cannot classify the resulting loss. It separates the mechanism from the ideological critique it is likely to be confused with, and takes up in detail the strongest existing objection, which is that recent power aware work in accessible computing has already displaced the account of adoption this paper is arguing against. It closes by building an instrument that measures the phenomenon from public artifacts, running a first pass, and reporting what that pass found.

One boundary is set at the outset. Whether accommodation repricing gives rise to obligations under disability discrimination law is a separate question with its own doctrinal apparatus, and it is not addressed here beyond acknowledging that it exists.

II. METHOD AND SCOPE

The paper has two parts and they answer to different standards.

The first is a conceptual argument about what two literatures contain, developed from published empirical work in accessible computing, disability studies, education, occupational therapy, and speech and language therapy, together with the published economics of switching costs. Its claims are falsifiable by counterexample: a taxonomy that already holds the category this paper says is missing would defeat it.

The second is an original documentary audit of vendor pricing artifacts, reported in Section IX. It exists because the conceptual argument makes a claim about supplier behavior, and a claim about behavior that rests only on other people's literature is an assertion. The audit is what makes it checkable.

A. Literature Selection

Four criteria. First, empirical studies of how disabled people use general purpose generative artificial intelligence, restricted to studies of deployed systems in ordinary use rather than elicitation studies of what disabled people would want. Second, the small body of work that names the financial dimension of inaccessibility as a concept rather than as an incidental finding. Third, the assistive technology abandonment literature, traced to its primary sources rather than to secondary syntheses, because the argument turns on what those taxonomies actually contain. Fourth, critical work on disability and technology capable of supplying the objections this argument must survive, which is treated in Section VII rather than confined to a related work section.

The third criterion did real work. A figure widely reported as an assistive technology abandonment rate of approximately 30 percent is frequently attributed to a 2004 monograph that is a synthesis rather than a study. The primary source is a survey of 227 assistive device users reporting an abandonment rate of 29.3 percent and identifying four contributing factors (Phillips and Zhao 1993). This paper uses the primary figure, and uses the four factors themselves as evidence rather than as background.

B. The Audit, and Why It Measures Suppliers

The audit collects published corporate artifacts, not accounts from people. That is a deliberate choice and not only a convenience.

The user side question is saturated. Recent work spans cross-disability focus groups, semi-structured interviews, autoethnography, workplace studies, diary studies, and video corpus analysis, and an additional small sample would add little. The supply side has never been measured at all. There is no time series anywhere in the accessibility literature showing what a given access capability cost, which tier it sat in, and whether its quantity was disclosed, tracked across years.

It is also the ethically cheaper instrument. Establishing that a vendor moved a term requires reading the vendor's own pages. Establishing what that did to a disabled person requires asking them to narrate a loss they could not prevent, which is worth doing only when they can be compensated for the account.

Analysis of published corporate web pages does not involve intervention or interaction with living

individuals and does not involve identifiable private information, and so falls outside the definition of human subjects research rather than within an exemption from it. Should a subsequent phase involve disabled participants, the design would be co-design with disabled people present in the framing, the priorities, and the evaluation, rather than consulted after the questions are fixed.

What the audit is not is a market analysis. The subject is a population, and the pricing record enters only as the stimulus that population is exposed to.

III. WHAT IS ESTABLISHED

A. Generative AI Is in Use as Accommodation

Disabled people use general purpose generative artificial intelligence to meet access needs. The evidence spans methods and communities.

A qualitative analysis of 25 videos from 25 autistic creators found the technology functioning as a digital coach, a communication assistant, and a conversational partner, with reported outcomes including harm reduction, savings of time and energy, and the facilitation of accommodations to environments that were themselves inaccessible (McNally, Wright, and Goldkind 2024). Interviews with 30 people with disabilities about their use of commercially available chatbots, of which 18 were analyzed, produced a mapping of distinct uses into nine broader categories, with writing and ideation the most common (Mullen, Xue, and Kudumu 2024). A subsequent study characterized uses and prompting strategies among neurodivergent individuals through interviews and a multi-week diary study (Xue et al. 2025). Work with autistic employees documented use for workplace and social communication (Jang, Moharana, and Carrington 2024). Autoethnographic work with neurodivergent sustained users documented the technology supporting code switching, emotional regulation, and information access, and described users reclaiming agency in systems that diminish their autonomy and self determination (Glazko et al. 2025).

Two features of this record matter here. The tools are mainstream commercial products, not devices dispensed through an assessment and funding pathway. And the same body of work has already begun to ask what happens over time. One of these studies states plainly that very little is known about how disabled people use these systems over an extended period (Glazko et al. 2025). That question is the one this paper takes up, from the supply side rather than the user side.

B. The Reliance Is Not Discretionary

Seven cross-disability focus groups with 20 participants found that while participants voiced substantial skepticism about the technology, many continued to use it out of necessity, and that where the technology is necessary for access or for safety, that necessity exposes disabled people to risks they would not otherwise accept (Johnson et al. 2026). This is the finding that converts a consumer product into an accommodation for analytic purposes. A person who continues to use a tool they distrust, because the alternative is losing access, is not exercising consumer choice in any sense the economics of consumer choice would recognize.

The same study describes participants navigating these tradeoffs through what accessible computing

has named a consequence calculus, a weighing of costs and consequences under conditions of constrained resources originally developed in work with chronically ill people (Mack, McDonnell, and Findlater 2022). That concept and the accessibility tax sit close together, and it is worth stating which one this paper extends. Consequence calculus describes a decision procedure a person runs. Accessibility tax describes a cost the person is charged. This paper extends the cost, and treats the decision procedure as the thing the cost is an input to. When the cost moves, the calculus is re-run on terms the person did not set.

C. The Vocabulary Exists

The financial dimension of inaccessibility has a name. The disability tax and the accessibility tax were distinguished as separate impositions, the first covering additional emotional and intellectual labor, the second covering the direct and indirect financial costs of navigating environments and technologies built on an assumption of nondisability (Olsen, Cork, and Anders 2022). The related term *crip tax* appears in academic use from 1999. The distinction has since been carried into generative artificial intelligence, where the accessibility tax was found to include financial costs, disclosure costs, and additional labor such as prompt engineering, and where pricing models were found to produce tiered accessibility in which financial resources determine the quality of access a disabled user receives (Johnson et al. 2026).

That last finding is the closest prior work to this paper's subject, and it is the foundation this paper builds on rather than a gap this paper fills.

D. What Is Not Established

What has not been described is the accessibility tax as a variable that moves.

Tiered accessibility is a snapshot: at this price, this quality of access. Accommodation repricing is the derivative: the tier boundary moves, or the price of the tier moves, or the quantity available inside the tier moves, after the user has organized their working life around the capability. The user did not change tiers. The tier changed around the user.

This distinction is not pedantic. A static tier is a barrier a person can assess before committing, and in principle can decline. A moving tier is a barrier that arrives after the decision has been made and after the switching costs have accumulated, at a moment when declining means giving up the accommodation. The two produce different harms, they call for different remedies, and only the first has been described.

IV. THE SUPPLY SIDE, STATED BRIEFLY

The economics here is not novel and is not this paper's contribution. It is stated as a premise because the mechanism it describes is the one the population in question is exposed to.

Suppliers facing customers with switching costs have a well characterized incentive structure: price low to build an installed base, then raise prices on customers who have become expensive to lose. The pattern is described in the switching cost literature as *bargain-then-ripoff* or *invest-then-harvest* (Klemperer 1987, 1995; Farrell and Klemperer 2007), and has been confirmed empirically in at least one large consumer market

where suppliers profited from consumer inertia (Ho, Hogan, and Scott Morton 2017).

Two refinements matter here.

First, suppliers passing through cost increases prefer to reduce quantity rather than post a price increase, because consumers are measurably less sensitive to quantity reduction than to price change, a preference that has been shown not to diminish with repeated exposure (Gourville and Koehler 2004; Çakır 2021; Kim 2023). Applied to a metered digital product, quantity reduction takes a distinctive form. The allowance can shrink without any published number changing, because in several cases no quantity term is published at all. Consumption is displayed to the user as a proportion of an allowance whose absolute size the vendor has not disclosed and reserves the right to alter. Under those conditions a reduction is not merely unnoticed. It is undetectable in principle from inside the product, which forecloses every user side response that depends on the user first perceiving that something has changed.

Second, the switching cost literature has been explicit about a limit in its own coverage: it models small buyers as passive, nonstrategic price takers, and does not address buyers able to act strategically (Lewis and Yildirim 2005). The population in this paper sits at the far end of that limitation. For a disabled user the switching cost is not a retraining expense that can be amortized. It is the loss of an accommodation.

The public record of the relevant period is consistent with the pattern. Between 2023 and 2026 major vendors of generative artificial intelligence products moved capabilities between subscription tiers, replaced flat allowances with metered credits, reduced allowances without notifying subscribers already inside paid annual terms, introduced usage allowances whose size was not published, and suspended new subscriptions on capacity grounds. Section IX reports an audit that converts part of that record into evidence.

V. THREE POSITIONS

A. *The Protections Attach to the Category*

Prescribed assistive technology arrives with an apparatus around it. There is an assessment, there is usually a funding pathway, there is an expectation of training, and there is some professional relationship responsible for the fit between person and device. Those obligations attach to the category of assistive technology, and they are the reason the category exists as an administrative object.

A commercial subscription has none of them. It has terms of service, a price, and a right of unilateral amendment.

The migration of function from the first category to the second has been described from inside occupational therapy. As mainstream products came to meet needs that had previously required specialist products, the result was a spectrum of products decoupled from access to independent advice and support services (Steel, Buchanan, and Layton 2017). That decoupling is the whole of this paper's structural claim, stated by practitioners in the field before generative artificial intelligence existed as a consumer product.

The same analysis observed that the purchase decision remains a leap of faith because these products are what economists call experience goods, whose value cannot be known until they are in use (Steel, Buchanan, and Layton 2017). That observation deserves more weight than it has received. An experience good whose

value is unknown before purchase, sold on a subscription with a reserved right of amendment, to a buyer for whom the product is an accommodation, is a set of conditions under which the buyer bears essentially all of the risk of the relationship.

A related warning was issued from disability studies in education: technology can create unexpected and under-critiqued forms of social exclusion for disabled people, and accessible options run behind whatever is being developed for mainstream markets (Foley and Ferri 2012). The present case is a variant in which the mainstream market did deliver the function, and the exclusion arrives through the commercial terms rather than through the delay.

B. The Third Position: The Procured Platform

There is a case the two-category framing does not reach, and it is the most exposed position of the three.

Accessibility functions are increasingly embedded in platforms procured by institutions rather than purchased by individuals. Speech to text, live captioning, automatic transcription, text to speech on documents, reading support, and summarization now appear as features inside learning management systems, productivity suites, classroom platforms, and education specific generative artificial intelligence offerings. A disabled student or employee who relies on one of those features is relying on an accommodation embedded in a product they did not select, do not pay for, and hold no contract with.

The adoption pattern in education is documented and it is fast. Emerging educational technologies including artificial intelligence, automated assessment, and online proctoring have been taken up rapidly across higher education, often without sufficient critical examination of how they affect disabled students and faculty, and analysis of learning management systems, artificial intelligence tutoring tools, and proctoring platforms shows that technologies which appear inclusive can simultaneously impose normative expectations and create new burdens for disabled people (Foley and Ferri 2026). A scoping review of 87 studies of generative artificial intelligence in higher education published between 2019 and 2026 identifies disability access as one of six equity dimensions running through the field, alongside the digital divide and socioeconomic stratification (Ni and Li 2026). Institutional acceptance and adoption of these models varies by institution, which itself produces inequality in access and use.

Against that background, consider what happens when an accessibility relevant feature inside such a platform is moved to a higher priced tier, metered, or withdrawn. The transaction that causes the loss occurs between the vendor and the institution. The disabled user is not a party to it. They have no notice right, no cancellation right, no refund, and no counterparty. They discover the change by finding that the feature is gone.

This position sharpens every element of the mechanism. The absence of leverage is total rather than merely practical, because there is no contract to leverage. The information asymmetry is complete, because the terms were never shown to the user at all; enterprise and institutional pricing is generally unpublished, so even the audit proposed in Section IX cannot see it from outside. And the substitution question becomes acute, because the institutional deployment of the platform is frequently the reason the user has no separate accommodation in place: the feature inside the platform was the accommodation, and its presence is why no

other arrangement was made.

It also makes clear that the phenomenon is not confined to standalone consumer subscriptions. Any product with an accessibility relevant capability and a pricing surface is capable of accommodation repricing, whether the disabled user is the purchaser, a dependent of the purchaser, or a person with no relationship to the purchase at all. Work in this area has already recommended that institutions conduct technology review and vendor evaluation with disability in view (Foley and Ferri 2026). The commercial terms under which an access capability is supplied, and the vendor's reserved right to alter them, belong on that evaluation list.

VI. ABANDONMENT WITHOUT A DECISION

A. *What the Taxonomies Count*

The assistive technology abandonment literature is the natural place to look for a category that fits accommodation repricing. It does not contain one.

The founding study surveyed 227 users of assistive devices across a range of disabilities and reported an abandonment rate of 29.3 percent, with four factors significantly related to non-use: the user's opinion was not taken into account during selection, the device was easily obtained, the device performed poorly, and the user's needs or priorities changed (Phillips and Zhao 1993). Every one of the four is a property of the artifact, of the process that matched the artifact to the person, or of the person's own trajectory. Not one describes an action taken by the supplier after the person began using the device.

The most substantial recent catalogue in augmentative and alternative communication reports eight contributing factors, synthesized from a survey of 43 people who use such systems and 68 professionals together with in-depth interviews with 18 users: ease of use, reliability, technical support, the voice and language of the device including intelligibility and personalization, the decision making process and lack of involvement in choosing the technology, service delivery and access to services, family perceptions and support together with staff training and partner competency, and communication rate (Waller 2019, reporting Judge and Townend 2013). A separate list in the same paper approaches abandonment from a design perspective and names physical access, slow communication rates, literacy demands, vocabulary organization, training requirements, and a design focus on needs based rather than extended communication.

Neither list contains price. A complete reading of that paper finds no discussion of cost, funding, affordability, or purchase anywhere in its body.

B. *The Objection, and the Answer*

The anticipated objection is that cost is already in the literature. It is, and the distinction matters.

Cost appears in the assistive technology literature as a barrier to acquisition. Devices are expensive, funding is limited, coverage is inconsistent, and people who need a device do not obtain one (Orellano-Colón, Mann, and Rivero-Méndez 2015). That is a real and well documented harm, and it is a property the device has at the moment of the acquisition decision. It is a fact about whether the person can get the thing.

Accommodation repricing is a fact about what happens after they have it. The person acquired the tool, adapted to it, built work and communication practices around it, and then the supplier moved a term. The abandonment that follows is not a decision the person made about fit, and it is not a change in their needs. It is a decision made elsewhere and delivered to them.

C. Why the Gap Persisted

The gap is not an oversight. It follows from the product form the literature grew up around. Abandonment research developed around durable devices with a one time acquisition and an ongoing service relationship. In that form the supplier's pricing decision is exhausted at the point of sale, so it appears in the taxonomy as an acquisition barrier and nowhere else. A subscription relocates the supplier's pricing decision to a point after adoption and makes it recurring, and the taxonomy has no place to put it.

That is a reason to extend the taxonomy rather than to fault it. The proposed extension is minimal: a category for discontinued use attributable to a change in supplier terms, distinct from user decision and from device failure, and recorded against the supplier rather than against the person or the product.

VII. OBJECTIONS

This argument has to survive four objections. Three come from existing literature and one is the strong form of a critique the paper is often going to be confused with. They are treated here rather than deferred, because the third in particular is the one most likely to be raised by a reader who knows this field well.

A. The Adoption-As-Fit Account Has Already Been Displaced

This is the strongest objection and it must be conceded in part.

Recent work in accessible computing argues directly that the choice to adopt or abandon assistive technology is often treated as a stable proxy for functional fit, that to adopt is read as confirming a good match between device features and individual needs and to abandon is read as signalling a poor one, and that this account is ill-equipped to explain the sociopolitical forces shaping assistive technology use in historically underserved communities. Drawing on an eight month ethnographic study, that work proposes a power aware account which recasts adoption as a situated negotiation of power rather than a transparent expression of fit, and introduces minor resistance to name the everyday tactics through which people lower the cost of access on their own terms. It concludes by reframing the goal of accessibility design from optimizing use to lowering the cost of choice (Hsueh, Van Dusen, and Caspi 2025). The same work observes that people will continue using an assistive technology even when it is unusable because they are structurally compelled to, and that non-use can be a legitimate and strategic response to structural tension rather than a failure of fit.

If the objection is that the adoption-as-fit account has already been challenged, the objection is correct and this paper does not claim otherwise. What remains is a difference in the kind of power at issue and in what follows from it.

The power described in that work is institutional and interpersonal: service systems, gatekeeping,

normative expectations about how disability should be performed, the cultures of the organizations through which devices are obtained. It is exercised by people and institutions in relationship with the disabled person, and the tactics of minor resistance are available precisely because those relationships are ongoing and negotiable at the margins.

The power described here is contractual and unilateral, exercised by a party with no relationship to the disabled person at all, through a term the person did not negotiate and in most cases never read. It is not exercised on the person; it is exercised on the product, and reaches the person through the product. There is no counterpart to minor resistance against a rate limit, and Section VII returns to why.

The two accounts are complementary and the second completes the first. That work asks accessibility design to lower the cost of choice. Accommodation repricing is a mechanism that raises the cost of choice, on a schedule the person does not control, after the choice has been made. Read together, the power aware account supplies the analytic move and this paper supplies a source of power it does not cover.

B. Disabled People Are Not Passive; They Adapt, Hack, and Refuse

Crip technoscience holds that disabled people are experts and designers of everyday life, that they harness technoscience for political action, and that they refuse to comply with demands to cure, fix, or eliminate disability, engaging in practices of critique, alteration, and reinvention of the material world (Hamraie and Fritsch 2019). Related work extends this to neurodivergent communication technologies (Rauchberg 2022) and to disabled and aging people's technologically mediated access to cultural life (Jones, Rice, and Lam 2021). Recent crip theory in artificial intelligence research reaches a compatible conclusion, challenging the assumption that accommodations inherently benefit disabled people and describing access as ongoing, friction filled work in which disabled people actively participate (Tang, Lin, and Li 2026).

The objection this line supports is that the present paper risks describing disabled people as passive recipients of a vendor decision, which is the very posture crip technoscience was written against.

The answer accepts the premise and denies the inference, and the denial is specific.

Every practice in the crip technoscience repertoire presupposes an artifact the user can get their hands on. Alteration, reinvention, tinkering, repair, hacking, workaround, and refusal all operate on something material and present. A wheelchair can be modified. A communication device can be reprogrammed, and its owner can decline to use it in the way the clinician intended. Those are real and consequential exercises of power over an object.

A metered subscription is not that kind of object. The constraint is enforced remotely, by contract and by server, on infrastructure the user does not possess and cannot inspect. There is nothing to open. Where the allowance is expressed as a proportion of an undisclosed quantity, there is not even a number to reason about. The one form of technology in the disabled person's life that is genuinely unalterable by them is the one that is not in their hands.

That is not a claim that disabled people are passive. It is a claim about a property of this specific product form, and it is the reason the crip technoscience repertoire does not reach it. Refusal remains available, and

non-use can be strategic (Hsueh, Van Dusen, and Caspi 2025). But refusal of an accommodation is a different act from refusal of a normative expectation, and the paper's argument is precisely that the cost of that refusal is being set by someone else and moved upward without notice.

C. This Is Just Technoableism

There is an existing and important critique of how technology is presented to and about disabled people. Technoableism names the belief that technology is the solution to disability, that disabled people are waiting to be repaired by technological advance, and that making the world accessible is a lower priority than developing the device (Shew 2023). Related work has named sociotechnical ableism and applied it to educational artificial intelligence systems including intelligent tutoring systems (Foley and Melese 2025), and an analytic instrument for educational technology synthesizes that critique with design justice and crip technoscience, arguing that technologies designed for disabled people often treat disability as a problem to fix rather than a valued aspect of human diversity, and thereby reproduce the inequities the institutions deploying them claim to address (Foley and Ferri 2026).

That critique is correct and this paper does not dispute it. It is also doing a different job, and conflating the two would weaken both.

The test is a counterfactual. Suppose a vendor's presentation of its product were beyond reproach: written by disabled people, free of every cure narrative, accurate about what the tool does and does not do, and explicit that the tool addresses an inaccessible environment rather than a deficient person. Under those conditions the harm described in this paper is unchanged, because the harm consists in the accommodation sitting behind a meter that someone else controls and can move. The framing critique does no work on the mechanism.

The relationship between them is sequential rather than competitive. Technoableism explains the arrival. If the governing assumption is that technology resolves disability and that structural accessibility is a secondary concern, then a market solution appears sufficient and the funded pathway is never built. That assumption is upstream of the situation this paper describes, and it is a reason the accommodation ended up as a commercial subscription instead of an entitlement with obligations attached. This paper takes up what happens next. Reading them in sequence assigns the framing critique its actual explanatory role, which is more than a paper that merely set it aside would do. Both directions of that reading are extensions of the framing critique rather than claims its authors make.

A narrower version of the same objection is that existing evaluative instruments for educational technology already cover this ground. They do not, and the reason is instructive. Those instruments evaluate what a technology does to disability as a category: whether it conserves disability as a valued form of human difference or normalizes it away, and whether it reflects disabled people's lived experience or an idealized narrative of unrestricted access (Foley and Ferri 2026). Those are questions about the design and the discourse. Accommodation repricing is a question about the commercial terms of supply, and a technology can score well on every design axis while being supplied on terms that permit its withdrawal next quarter. The two evaluations are orthogonal and both are needed.

D. A New Name Is Not Warranted

For the same reason this paper declines to coin a new category of ableism. A new term must do work that existing terms do not. Here both ends of the mechanism are already named and each has a literature a reader can check: the accessibility tax on the cost side, and the invest-then-harvest pattern on the supply side. Placing an established term at each end is stronger than a neologism at the middle, because both ends are independently verifiable.

Accommodation repricing is offered as a label for an event, not as a theory of a new form of oppression. It claims that a specific thing happens and invites the reader to check whether it does. It is meant to be checkable rather than defended.

VIII. WHO THE MECHANISM REACHES

The population exposed to accommodation repricing is wider than the population visible in studies of paying subscribers, and four subclasses deserve naming.

Disabled people who were priced out before the change occurred appear in no study of users, because they are not users. Any measurement of harm taken from the subscriber base understates the harm by exactly the number of people the price has already excluded.

Disabled people for whom the subscription is a larger share of income experience the same nominal increase as a different event. Where disability is associated with lower employment and lower income, an increase that is an inconvenience at one income is a cancellation at another, and the distribution of harm inside the class is uneven in a way an average will conceal.

Disabled people whose access capability sits inside an institutionally procured platform, as described in Section V, hold no contract and therefore have no notice, no cancellation, and no remedy in the ordinary consumer sense.

The fourth subclass is the one where the standard remedy and the harm collide. For disabled people whose impairment affects the ability to track a recurring charge, register a change of terms, or act on a notice within a window, the usual proposed fix is precisely the thing that fails. Notice obligations are the conventional answer to unilateral amendment. Notice delivered by email, buried in a change log, or presented as a modal dialogue during a task is not notice that reaches a person whose executive function, attention, memory, or sensory access is the site of their disability. A remedy that consists of telling people what is about to happen assumes a reader the design of the remedy has not accounted for.

That subclass should be treated as a design requirement rather than an edge case. Any remedy proposed for accommodation repricing that depends on the user noticing something must specify how it reaches the people least able to notice, or it is not a remedy for them. The point compounds with the disclosure problem in Section IV: where the quantity term is unpublished, there is nothing for even an attentive user to notice.

IX. THE AUDIT

The phenomenon is measurable without recruiting a single participant, because the artifacts that constitute it are published.

The instrument is a longitudinal audit of vendor pricing artifacts across the 2023 to 2026 window: pricing pages, plan comparison tables, published allowance and rate limit documentation, terms of service, help documentation describing accessibility relevant capabilities, and changelogs. Historical states are recovered from public web archives, so every observation carries a dated capture with a citable address.

The unit of analysis is the vendor and capture date. The coding scheme records, for each accessibility relevant capability observed: what the capability is, expressed in access terms rather than marketing terms, such as speech to text, live captioning, image description, document text to speech, long context retention, or summarization; which tier the capability sits in and at what price; whether the quantity term is published as an absolute number, published as a relative proportion, not published at all, or expressed in a unit that is not commensurable with the unit used in the preceding capture; whether a change was announced in advance and through what surface; whether the change reached subscribers already inside a paid term; whether any accessibility specific concession exists, such as reduced pricing on disability grounds or a partnership arrangement with a disability organization; and whether the capability remains available in any tier or has left the consumer product entirely.

The audit covers two vendor classes rather than one. The first is consumer generative artificial intelligence subscriptions. The second is institutionally procured platforms carrying accessibility relevant capabilities, including learning management systems, classroom platforms, productivity suites, and education specific artificial intelligence offerings. The second class is where the procured platform position in Section V lives, and its published surfaces, while thinner, still include feature comparison pages, education plan tables, and support documentation stating which capabilities belong to which edition.

Three outputs follow. A time series per capability showing tier position, price, and disclosure state, which is the empirical object neither existing literature has. A disclosure typology distinguishing published quantity terms from unpublished ones, which operationalizes the undetectability point in Section IV. And a notice typology recording which changes were announced, to whom, through what channel, and how far in advance, which is the input any serious remedy proposal will need.

A. Preliminary Findings

A first pass has been run. Sixteen surfaces across both vendor classes were indexed against the Internet Archive for the 2023 to 2026 window, returning archived states for 15 of them; one education vendor produced no captures and is reported as a gap rather than as an absence of change. One consumer surface, the plan comparison page for a widely used commercial coding assistant, was sampled at quarterly intervals and coded in full across 12 consecutive quarters from the fourth quarter of 2023 to the third quarter of 2026. It is offered here as a demonstration that the instrument registers the phenomenon, not as a general result. The capture index, the coding scheme, and the coded series are deposited in full (Gilly 2026).

The headline price of the individual paid tier was ten dollars per month at the first capture and ten dollars per month at the last. Across three years it never moved. A study designed to detect price increases would have found nothing.

What moved was quantity, and then the terms in which quantity was expressed. In the second quarter of 2024 the paid tiers carried an explicit statement of unlimited messages and interactions. By the third quarter of 2025 a metered unit had been introduced alongside that claim: the unlimited guarantee was narrowed to a single named model, and access to the current models was capped at 300 units per month on the same ten dollar tier, with additional units sold at four cents each. In the third quarter of 2026 the metered unit was replaced altogether by a balance denominated in dollars, and the same tier carried fifteen dollars of monthly credit.

That last change is the one the coding scheme was extended to capture. A user cannot determine from the two surfaces whether fifteen dollars of credit is more or less than 300 units of the previous kind, because the vendor did not publish a conversion and the unit of account changed. The comparison a person would need to make in order to notice a reduction is not available to them, and it is not available because of how the change was expressed rather than because the person failed to look.

A second observation concerns where the numbers live. At the final capture the plan comparison table describes the free tier's allowance as limited chat and agent usage, with no figure. The figure has not been deleted; a complete reading of the page finds it surviving in a footnote. This is recorded as a partial rather than a total withdrawal of disclosure, and the distinction matters. The claim is that a quantity moved off the surface where a person choosing between plans would encounter it, not that the vendor concealed it.

The sequence is what the argument predicted. A supplier holding customers with switching costs left the posted price untouched and adjusted the quantity behind it, which is the behavior the consumer economics literature describes, and then altered the unit in which that quantity was reported. For a general customer this is a pricing question. For a customer whose use of the product is an accommodation, each of those adjustments lands on a capability they have organized their working life around, and none of them was announced as a price increase because none of them was one.

The instrument has clear limits. It measures the stimulus and not the response, so it cannot establish how many disabled people lost access or what it cost them; it establishes only that the terms moved and how. It is restricted to vendors with archived public pricing surfaces, which biases toward larger firms in wealthier markets. Enterprise and institutional pricing is largely unpublished, which is a particular limitation for the procured platform position, where the operative transaction is precisely the one not on a public page. Archival coverage is uneven, so absence of a capture is not evidence of absence of a change. And vendors move. Two of the surfaces indexed for this study ended abruptly in September 2025 and resumed within a week at a different domain, so an auditor who does not check for relocation will record a discontinued product where there was only a redirect.

X. THE DESIGN SPACE

Four directions follow from the mechanism. They are set out as a space rather than a recommendation, and no particular product or provider is endorsed.

Local execution. Where an access task does not require frontier capability, running it on the user's own hardware removes it from a meter entirely. The cost becomes a one time purchase and electricity, and the supplier loses the ability to reprice a capability the user already holds. This is also the one direction that restores the property Section VII.B identified as missing, because a capability running on the user's own device is an artifact the user can get their hands on.

Notice and change obligations. Advance notice of price, tier, and quantity changes, published absolute quantity terms rather than proportions of undisclosed allowances, and protection of subscribers inside a paid term. Subject to the reaching problem in Section VIII, which any such obligation must address rather than assume away.

Portability. The accumulated context that makes a tool useful as an accommodation, including stored preferences, custom instructions, and history, is itself a switching cost. Making it exportable and importable reduces the cost of leaving and therefore reduces the supplier's ability to harvest.

Coverage and procurement. Extending assistive technology funding pathways to cover mainstream subscriptions where they serve an assistive function would restore the obligations that the migration described in Section V stripped away. In the procured platform position the corresponding lever sits with the institution rather than the individual: contractual continuity terms for accessibility relevant capabilities, notice to affected users rather than to the purchasing office alone, and inclusion of supply terms in the vendor evaluation that disability aware technology review already recommends (Foley and Ferri 2026).

XI. LIMITATIONS

The core claim is conceptual and rests on an argument about what two literatures do and do not contain. That claim is falsifiable in an obvious way: a taxonomy of assistive technology abandonment that includes supplier initiated change in terms as a category would defeat it. The two examined here do not, but the literature is large and the claim is stated about those two.

No causal claim is made. This paper does not establish that any disabled person has lost an accommodation because of a pricing change. It establishes that the conditions for that loss exist, that the population exposed to it has been documented as dependent, and that no existing taxonomy would record the loss correctly if it occurred. Establishing incidence requires work not done here.

The audit measures suppliers, not people, and is silent on the magnitude of harm. It is weakest exactly where the argument is strongest, because the procured platform position turns on transactions that are not published. The reported pass is also narrow: 16 surfaces indexed, one coded in full. A single coded series demonstrates that the instrument registers the phenomenon and cannot establish how common the phenomenon is. Nothing in Section IX should be read as a prevalence claim.

The pricing record used as background is drawn from a period of unusual volatility in one industry. Whether the pattern generalizes beyond it is an open question, though the switching cost literature suggests the incentive is not specific to this industry.

Finally, accommodation repricing as offered here is a label, not a construct with established measurement properties. It has not been validated, and the term should be judged on whether it makes a real phenomenon easier to see and to check.

XII. FUTURE WORK

Two items, each answering a specific limitation above.

Extending the audit answers the absence of any time series on the supply side. The reported pass covers one coded surface; coding the remaining indexed surfaces across both vendor classes would convert a demonstration into a prevalence estimate, and requires no participants.

A prospective study of disabled users through an announced pricing change would answer the causal gap. The design that would work is a cohort recruited before a scheduled change takes effect, with access practices documented in advance, followed through the change, in co-design with disabled researchers and with participants compensated. It should be attempted only where compensation is funded, because asking disabled people to narrate a cost increase they cannot escape, without paying them for the account, reproduces the imposition the study is meant to document.

XIII. CONCLUSION

The accessible computing literature has established that disabled people use general purpose generative artificial intelligence to meet access needs, that the reliance is frequently necessary rather than optional, and that what a person can pay determines the quality of access they receive. This paper has argued that the last of those findings is stated as a level and needs to be stated as a change.

When the accommodation is a commercial subscription, the price, the metered quantity, and the capability are all terms the supplier may move after the user has reorganized around them. The protections that would attach to the same function delivered as prescribed assistive technology do not travel with the function into the commercial category. The abandonment literature, which exists to explain why access tools fall out of use, attributes that loss to the artifact, to the match, or to the person, and has no category for a loss produced by a supplier decision made after adoption. Where the capability is embedded in a platform procured by an institution, the disabled user is not even a party to the transaction that removes it, and the transaction itself is not published.

Naming the event is the modest part of the contribution. The substantive part is the claim that the loss it produces is currently unrecordable: it is not visible as an accessibility tax, because that concept describes a standing cost; and it is not visible as abandonment, because that concept describes a user decision. A harm that no existing instrument can register is a harm that will be counted as nothing, and this one is arriving on a schedule set by someone other than the people it lands on.

DATA AVAILABILITY

The audit reported in Section IX is reproducible from public sources. The target list, the harvest and sampling code, the capture index, the coding scheme, and the dated public archive address of every coded capture are available at <https://github.com/princesshobbes/accommodation-repricing-audit> (Gilly 2026). Code is released under the MIT licence and data and documentation under CC BY 4.0. An archived release under a persistent identifier will accompany publication.

Archived vendor pages are not redistributed, and remain the property of their respective owners. Each coded observation cites the public archive address from which it was read, so any reader can retrieve the same page and check the coding against the source rather than against a copy hosted by the author. Coding errors reported against a row number and an archive address will be corrected in the repository record.

ACKNOWLEDGMENTS

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic’s Claude as assistive technology for a diagnosed neurodevelopmental condition. All intellectual contributions, research questions, theoretical arguments, methodological decisions, and conclusions are solely the author’s own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: t.gilly@ai-literacy-labs.org.

REFERENCES

- Çakır, M. (2021). Retail pass-through of package downsizing. *Agribusiness*, 38(2), 259–278. <https://doi.org/10.1002/agr.21724>
- Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong & R. Porter (Eds.), *Handbook of Industrial Organization* (Vol. 3, pp. 1967–2072). Elsevier.
- Foley, A., & Ferri, B. A. (2012). Technology for people, not disabilities: Ensuring access and inclusion. *Journal of Research in Special Educational Needs*, 12(4), 192–200. <https://doi.org/10.1111/j.1471-3802.2011.01230.x>
- Foley, A., & Ferri, B. A. (2026). DISCO-Tech: A framework for resisting ableism in emerging technologies. *Teachers College Record*, 128(3), 166–194. <https://doi.org/10.1177/01614681261444567>
- Foley, A., & Melese, F. (2025). Disabling AI: Power, exclusion, and disability. *British Journal of Sociology of Education*. <https://doi.org/10.1080/01425692.2025.2519482>
- Glazko, K., Lewis, A., Kosa, B., et al. (2025). Autoethnographic insights from neurodivergent GAI “power users.” In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3706598.3713670>
- Gourville, J. T., & Koehler, J. J. (2004). *Downsizing price increases: A greater sensitivity to price than quantity in consumer markets* (Working Paper). Harvard Business School. <https://doi.org/10.2139/ssrn.559482>
- Kim, I. K. (2023). Consumers’ preference for downsizing over package price increases. *Journal of Economics & Management Strategy*, 33(1), 25–52. <https://doi.org/10.1111/jems.12548>
- Gilly, T. (2026). *Accommodation repricing audit* [Data set and code]. Real Safety AI Foundation. <https://github.com/princesshobbes/accommodation-repricing-audit>
- Hamraie, A., & Fritsch, K. (2019). Crip technoscience manifesto. *Catalyst: Feminism, Theory, Technoscience*, 5(1), 1–33. <https://doi.org/10.28968/cftt.v5i1.29607>
- Ho, K., Hogan, J., & Scott Morton, F. (2017). The impact of consumer inattention on insurer pricing in the Medicare Part D program. *RAND Journal of Economics*, 48(4), 877–905.
- Hsueh, S., Van Dusen, D., & Caspi, A. (2025). Minor resistance: The everyday politics and power dynamics of assistive technology adoption. In *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility*. ACM. <https://doi.org/10.1145/3663547.3746465>
- Jang, J., Moharana, S., & Carrington, P. (2024). “It’s the only thing I can trust”: Envisioning large language model use by autistic workers for communication assistance. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM. <https://doi.org/10.1145/3613904.3642894>
- Johnson, J., Lewis, A., Mankoff, J., & Banner, O. (2026). “I don’t trust it, but I use it”: Navigating trust, privacy, and identity in disabled people’s use of generative AI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (pp. 1–17). ACM. <https://doi.org/10.1145/3772318.3790652>
- Jones, C. T., Rice, C., & Lam, M. (2021). Toward TechnoAccess: A narrative review of disabled and aging experiences of using technology to access the arts. *Technology in Society*, 65, 101537. <https://doi.org/10.1016/j.techsoc.2021.101537>
- Judge, S., & Townend, G. (2013). Perceptions of the design of voice output communication aids. *International Journal of Language and Communication Disorders*, 48(4), 366–381.
- Klemperer, P. (1987). Markets with consumer switching costs. *Quarterly Journal of Economics*, 102(2), 375–394.
- Klemperer, P. (1995). Competition when consumers have switching costs. *Review of Economic Studies*, 62(4), 515–539.
- Mullen, K., Xue, W., & Kudumu, M. (2024). “I’m treating it kind of like a diary”: Characterizing how users with disabilities use AI chatbots. In *Proceedings of the 26th International ACM SIGACCESS Conference on*

- Computers and Accessibility (Article 133, 7 pp.). ACM. <https://doi.org/10.1145/3663548.3688549>
- Lewis, T., & Yildirim, H. (2005). Managing switching costs in multiperiod procurements with strategic buyers. *International Economic Review*, 46(4), 1233–1269.
- Mack, K., Martinez, J. J., & Lewis, A. (2025). Modeling accessibility: Characterizing what we mean by “accessible.” In *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility*. ACM. <https://doi.org/10.1145/3663547.3746344>
- Mack, K., McDonnell, E. J., & Findlater, L. (2022). Chronically under-addressed: Considerations for HCI accessibility practice with chronically ill people. In *Proceedings of the International ACM SIGACCESS Conference on Computers and Accessibility*. ACM. <https://doi.org/10.1145/3517428.3544803>
- McNally, K., Wright, K., & Goldkind, L. (2024). Disability expertise and large language models: A qualitative study of autistic TikTok creators’ use of ChatGPT. *Social Media + Society*, 10(3). <https://doi.org/10.1177/20563051241279549>
- Ni, W., & Li, J. C. (2026). Mapping the impact of generative AI in higher education: A scoping review of psychological and equity dimensions. *Frontiers in Psychology*, 17. <https://doi.org/10.3389/fpsyg.2026.1856854>
- Olsen, S. H., Cork, S. J., & Anders, P. (2022). The disability tax and the accessibility tax. *Including Disability*, (1), 51–86. <https://doi.org/10.51357/id.vi1.170>
- Orellano-Colón, E. M., Mann, W. C., & Rivero-Méndez, M. (2015). Hispanic older adults’ perceptions of personal, contextual and technology-related barriers for using assistive technology devices. *Journal of Racial and Ethnic Health Disparities*, 3(4), 676–686.
- Phillips, B., & Zhao, H. (1993). Predictors of assistive technology abandonment. *Assistive Technology*, 5(1), 36–45. <https://doi.org/10.1080/10400435.1993.10132205>
- Rauchberg, J. S. (2022). Imagining a neuroqueer technoscience. *Studies in Social Justice*, 16(2), 370–388. <https://doi.org/10.26522/ssj.v16i1.1111>
- Shew, A. (2023). *Against technoableism: Rethinking who needs improvement*. W. W. Norton.
- Steel, E., Buchanan, R., & Layton, N. (2017). Currency and competence of occupational therapists and consumers with rapidly changing technology. *Occupational Therapy International*, 2017, 1–5. <https://doi.org/10.1155/2017/5612843>
- Tang, X., Lin, T., & Li, J. (2026). Crippling AI: Reimagining AI through lived disability experiences. ACM. <https://doi.org/10.1145/3805689.3806744>
- Xue, W., Kudumu, M., Sriram, S., Mullen, K., Boyd, A., & Gadiraju, V. (2025). Characterizing uses and prompting strategies of LLM-based chatbots among neurodivergent individuals. In *Proceedings of the 27th International ACM SIGACCESS Conference on Computers and Accessibility*. ACM. <https://doi.org/10.1145/3663547.3759721>
- Waller, A. (2019). Telling tales: Unlocking the potential of AAC technologies. *International Journal of Language and Communication Disorders*, 54(2), 159–169. <https://doi.org/10.1111/1460-6984.12449>