

Towards Designing for Resilience: Community-Centered Deployment of an AI Business Planning Tool in a Pittsburgh Small Business Center

Quentin Romero Lauro
School of Computing and Information
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
qtr2@pitt.edu

Aakash Gautam
School of Computing and Information
University of Pittsburgh
Pittsburgh, Pennsylvania, USA
aakash@pitt.edu

Yasmine Kotturi
Human-Centered Computing
University of Maryland, Baltimore
County
Baltimore, Maryland, USA
kotturi@umbc.edu



Figure 1: We introduced BizCHAT, an AI-powered business planning tool for small business owners, and deployed it through co-designed, community-centered workshops at a feminist makerspace in Pittsburgh. This figure shows P5, a custom jewelry and apparel designer, using BizCHAT to create her first business plan and gain entry to a grant-backed fashion week program.

Abstract

Entrepreneurs in resource-constrained communities often lack time and support to translate ideas into actionable business plans. While generative AI promises assistance, most systems assume high digital literacy and overlook community infrastructures that shape adoption. We report on the community-centered design and deployment of BizChat, an AI-powered business planning tool, introduced across four workshops at a feminist makerspace in Pittsburgh. Through log data ($N=30$) and interviews ($N=10$), we examine how entrepreneurs build resilience through collective AI literacy development—encompassing adoption, adaptation, and refusal of AI. Our findings reveal that while BizChat lowered barriers to accessing capital by translating ideas into "business language," this ease raised questions about whether instant AI outputs undermine sensemaking essential to planning. We show how peer support

helped entrepreneurs navigate this tension. We contribute design implications, including productive friction, communal scaffolds, and co-optability, for strengthening resilience amid technological change.

CCS Concepts

• Human-centered computing → Human computer interaction (HCI).

Keywords

Small Business, Entrepreneurship, Business Planning, Generative AI, Resilience

ACM Reference Format:

Quentin Romero Lauro, Aakash Gautam, and Yasmine Kotturi. 2026. Towards Designing for Resilience: Community-Centered Deployment of an AI Business Planning Tool in a Pittsburgh Small Business Center. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3772318.3791654>



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/26/04

<https://doi.org/10.1145/3772318.3791654>

1 Introduction

Small business owners are uniquely positioned to benefit from generative AI, particularly for automating tedious back-office tasks where time and resource constraints often hinder growth. Prior work suggests that large language models (LLMs) can improve entrepreneurs' bottom line by streamlining workflows, soliciting and implementing feedback, and unlocking time for long-term planning [69]. These systems can produce high-quality drafts for written content, support decision-making, and leverage broad knowledge bases to adapt to generalist tasks.

However, integrating generative AI into entrepreneurial workflows is far from straightforward. While commercial systems, such as ChatGPT or Gemini, present themselves as simple to use [67], their successful adoption requires a range of "hidden" operational skills [45]. Entrepreneurs must navigate browser literacy, password management, cloud storage, and other foundational practices that extend well beyond prompt engineering [48]. Even when entrepreneurs successfully generate AI outputs, a gap remains in translating generic AI-generated text into contextualized, business-specific content that reflects their actual operations, customers, and goals [48].

These barriers compound for entrepreneurs in resource-constrained settings, where limited access to social capital, infrastructure, and time for digital upskilling creates widening disparities in AI adoption and outcomes [23, 39, 41, 50]. For many small business owners, accessing capital requires business plans—formal documents that articulate goals and strategies for growth [89]—but writing these plans demands both business planning expertise and comfort with "the language of business" that funders expect. This double barrier of needing to both understand what a business plan should contain and possess the writing fluency to articulate it leaves many entrepreneurs unable to compete for funding, even when their business ideas are sound. While users with higher technical fluency can leverage generative AI to accelerate such tasks, others face widening disparities in use and outcome [18, 34]. Over-reliance on generative AI for tasks ill-suited to automation can further exacerbate these divides, reinforcing "rich-get-richer" dynamics in technology adoption [69].

In this work, we examine how small business owners navigate these compounding challenges through two concepts. First, we draw on resilience—patterns of positive adaptation in the context of adversity [77]—as a collective capacity that emerges from shared resources, mutual support, and local infrastructures [78]. We operationalize resilience as the shift from uncertainty and hesitation in the face of technological change ("unknown unknowns") to collectively arriving at possibilities of action ("known unknowns") [87, 94]. Second, we use AI literacy to refer to both operational competence with AI systems and the ability to adopt, adapt, or refuse AI tools based on one's context [13, 57]. These concepts are closely intertwined, as AI literacy enables resilience by transforming opaque, threatening technology into something entrepreneurs can evaluate strategically, and resilience emerges when communities collectively develop that literacy rather than navigating technological change alone. Building on this, we argue that for entrepreneurs in 2025, business resilience and AI literacy are inseparable: the pressure to compete for capital, maintain professional credibility, and stay

current with evolving tools means that navigating AI is itself a core business challenge.

We partnered with Prototype PGH, a feminist makerspace and entrepreneurial hub in Pittsburgh, Pennsylvania dedicated to racial and gender equity in technology. Building on formative research [48], we co-developed and deployed BizCHAT¹, an LLM-powered web application that supports product- and service-based entrepreneurs in writing business plans. Unlike commercial offerings such as Microsoft 365 Copilot, which rely on users' ability to craft sophisticated prompts [61], BizCHAT scaffolds both business planning by centering the entrepreneur's existing knowledge, and adapting the technology to the entrepreneur's context, rather than requiring the entrepreneur to adapt to the technology.

BizCHAT's design operationalizes the two concepts—resilience and AI literacy—through three corresponding design goals. First, we adopt a low-floor-high-ceiling [76] approach to aid users with limited digital skills or no prior knowledge of generative AI while offering advanced features for entrepreneurs experienced with tools like ChatGPT, ensuring that the tool supports AI literacy development across the full spectrum from adoption through adaptation to refusal. To facilitate skill-building for time-constrained entrepreneurs [9], our second design consideration is just-in-time learning [38, 82], which scaffolds AI literacy incrementally as entrepreneurs work toward their business goals rather than requiring upfront training. Finally, we explicitly contextualize the introduction of generative AI within the users' existing knowledge and goals (i.e., their business), rather than highlighting the novel technology [10, 30], supporting resilience by grounding AI adoption in entrepreneurs' concrete expertise and immediate needs.

From January through August 2025, we integrated BizCHAT into a four-session workshop series at Prototype. Workshops were framed as "Resilience and Uncertainty Planning" sessions focused on business planning, with AI introduced as one tool among many for drafting plans. This framing was deliberate: by positioning business planning—not AI—as the focal point, we created space for entrepreneurs to engage with technology on their own terms while surfacing questions about adoption, trust, and control. In total, 30 entrepreneurs used BizChat, including workshop participants (N=21) and organic users who adopted through word-of-mouth (N=9). We collected system log data and conducted 10 semi-structured follow-up interviews to understand how entrepreneurs navigated AI-assisted business planning within existing support infrastructures.

Rather than evaluating BizCHAT in isolation, we investigate how its design and deployment interact with existing practices of support, learning, and adaptation, and how participants navigated the combined challenges of business planning, AI adoption, and capital access. From this perspective, we ask the following research questions:

- RQ1. How should AI systems be designed to lower existing barriers of adoption for entrepreneurs in resource-constrained contexts?

¹BizCHAT can be freely used at: <http://bizchat-io.vercel.app/>. We named our tool BizChat before Microsoft Copilot Business Chat was released—since the tool is increasingly established within the community of entrepreneurs we serve, we decided not to rename.

- RQ2. How can existing community infrastructures mediate the development of AI literacy—including adoption, adaptation, and refusal of AI tools—for entrepreneurship?
- RQ3. In what ways does resilience emerge when AI tools are deployed within existing communities of entrepreneurial support?

By designing for low-floor-high-ceiling accessibility (e.g., voice-to-text for entrepreneurs with limited typing proficiency), just-in-time learning (e.g., scaffolded onboarding questions), and contextualized AI introduction (e.g., extracting information from entrepreneurs' existing websites), BizCHAT aimed to lower barriers to adoption [RQ1]. Our findings reveal that while BizCHAT reduced barriers to accessing capital by scaffolding the drafting process and "leveling the language of business," this very ease introduced tensions between relying on instant outputs versus engaging in the deeper sensemaking and reflection essential to effective planning [RQ1]. We show how peer support helped participants overcome technology gaps—such as contextualizing generic AI text for their specific customers, judging when outputs were accurate versus plausible-sounding, and translating business jargon into their authentic voice—by collectively developing strategies for editing, refining, and assessing output quality [RQ2]. Further, while BizCHAT seeded peer-led infrastructures for business planning and AI adoption, participants debated whether AI tools should be introduced by trusted community members or external experts, revealing tensions around credibility, expertise, and the appropriate pace of technology adoption [RQ2]. Participants demonstrated resilience against technological exclusion and capital barriers by collectively transforming their initial fear, uncertainty, and confusion about AI—an "unknown unknown"—into concrete strategies for adopting, adapting, and refusing AI [RQ3].

Taken together, this paper contributes: (1) empirical insights from the log data and community-based deployment of BizCHAT with 30 entrepreneurs; (2) a methodological framing of resilience in technology adoption, embedding it into workshops and accountability interviews; and (3) practical design and deployment implications for introducing AI systems in resource-constrained communities, including the value of productive friction, communal scaffolds for AI literacy, and designing for co-optability.

2 Related Work

2.1 The Hustle Economy Meets AI

The rapid rise of the "hustle" economy—characterized by growing numbers of freelance, solo-entrepreneurial ventures and side-hustle culture—reflects shifting work norms where individuals increasingly pursue flexible, short-term, and independent income streams across diverse sectors [16]. Today, nearly half of U.S. small businesses with paid staff operate at an extremely modest scale, with 49% (or 2.9 million employer firms) employing just one to four workers, showing that much of the small business landscape resembles solopreneurship or micro-enterprise rather than larger employer firms [54]. Technology has accelerated this shift by reducing startup costs, expanding customer reach through e-commerce and mobile platforms, and lowering barriers to entry, widely touted as a "democratizing force" for innovation, economic mobility, and social change [91]. While positive narratives of "hustle culture" cast

entrepreneurial pursuits as aspirational pathways to autonomy and self-realization [19], HCI and business researchers caution that this framing obscures the economic precarity underlying many ventures, where individuals pursue entrepreneurship less out of choice and more out of necessity [53, 63]. Scholars describe this "necessity-driven entrepreneurship" as self-employment undertaken in response to limited job opportunities, systemic employment barriers, or unstable economic conditions — a pathway for survival rather than aspiration [41, 63].

Against this backdrop, AI raises pressing questions for HCI scholars: will these tools expand opportunities for necessity-driven entrepreneurs, or further entrench divides in access and capacity? On the one hand, a growing ecosystem of AI tools promises to help small-scale business owners tackle their endless to-do lists through automation and outsourcing (e.g., *Storyteq* for automated ads creation, *Clay* for automated leads generation, business planning platforms like *LivePlan* and *Upmetrics*), indicating promise towards democratizing entrepreneurship among those with internet access. But in practice, successful adoption requires what scholars call a "hidden curriculum" of digital work [45] — a "laundry list of operational skills" entrepreneurs must acquire in order to effectively integrate AI into small business, including basic browser literacy, file-type conversions, password management, and clipboard know-how, among others [48]. But even after acquiring these accessory skills, entrepreneurs from lean-economies [23] face deeper obstacles: resource constraints and historical infrastructural barriers to technology adoption [9, 39], ultimately perpetuating a "rich gets richer" dynamic in AI adoption [69]. Addressing such divides requires moving beyond individual solutions to foreground community and structural barriers [39]. We extend this scholarship by deploying a system addressing these tensions and uneven conditions across design, development, and adoption.

2.2 Resilience Through Collective Capacity Building

Amid rapid technological change, people must adapt their lifestyles, workflows, and everyday practices. Resilience is widely recognized as a critical capacity for navigating these shifts [55, 60, 77]. Scholars in psychology and human development define resilience as "patterns of positive adaptation in the context of past or present adversity, one class of phenomena observed in human lives" [60]. In practice, scholars have detailed various strategies people use to exercise resilience, such as developing new routines to develop emotional stability in the face of loss or trauma [15], coping by re-framing challenges as manageable amid chronic stress [79], or by cultivating self-regulation and persistence to maintain academic performance amid economic hardship [58]. However, resilience is not an individual capacity alone [66, 72, 83, 88]. It is also embedded in and functions as a collective capacity. Resilience is related to but distinct from social support and collective learning. Social support provides resources; collective learning builds skills; but resilience describes how communities transform uncertainty and threat into actionable knowledge and sustained practice [60, 77]. This distinction matters here, where entrepreneurs face not just skills or resource gaps, but widespread uncertainty about engaging with

rapidly evolving, potentially threatening technology. Resilience captures the adaptive capacity to navigate this uncertainty — turning “unknown unknowns” into “known unknowns” that can be named, discussed, and acted on collectively [87, 94].

HCI scholarship demonstrates how this transformation can occur in practice. Scholars have shown how people draw on social and material resources within their communities to adapt when facing adversity [92]. For example, scholarship on assets-based design advocates identifying and building on community assets, such as collective knowledge, shared resources, and peer support, as a way to enable communities to realize goals that matter to them (e.g., [30, 73, 95]). Collective capacity building can scaffold resilience by allowing people to pool resources and knowledge [39, 50, 92]. In entrepreneurial contexts, for instance, Hui et al. [39] describe “Community Collectives,” where small groups of necessity-driven entrepreneurs support one another in overcoming barriers to digital platform use through low-tech coordination and regular gatherings. Similarly, Kotturi et al. [50] examine a “Tech Help Desk” model in which a small business support program provided long-term, one-on-one technical assistance. Through collaboratively addressing the “long tail” of everyday computing challenges, entrepreneurs solved immediate problems, and also developed strategic skills (e.g., goal setting) and operational skills (e.g., implementation) for adapting to technological change [50]. Other work highlights how community organizations act as trusted intermediaries that broker technical knowledge, translate platform norms, and provide ongoing, place-based support [22, 36]. The scholarship highlights the often hidden, varied, but critical non-technological infrastructures that are necessary in supporting communities, such as entrepreneurs in lean economies [23, 24].

While prior scholarship shows how resilience can be scaffolded in response to economic and digital inequities, AI introduces a new layer of fragility. Recently, Glassman et al. [31] propose AI-resilient interfaces designed to help users notice, judge, and recover from AI-generated errors. These practices are elements of building resilience and can be leveraged to facilitate collective capacity building. For example, when noticing and recovering from errors are facilitated through group-based sensemaking and leveraging domain knowledge that participants already have, it can build collective capacity and strengthen resilience. Building on these insights, we position technological capacities within the community and explore how we can support resilience building through a network of support—where our design, research, and deployment are aligned to scaffold community interactions and collective engagement.

An equally important part of resilience in this setting is the capacity for refusal. Refusal refers to the choice not to use a technology, or to use it selectively, due to misalignment with one’s goals, values, or constraints [11, 27, 80, 81]. For instance, Kotturi et al. [48] demonstrate entrepreneurs’ hesitation around using AI for grant writing or marketing copy due to fears of inaccuracy or legal ambiguity. Designing for resilience, then insists on designing for informed refusal where users are supported to evaluate when and how AI fits their needs.

Building on this scholarship, we frame resilience in our work as the community’s capacity to adopt, adapt, and refuse AI. We operationalize this framing through BizCHAT, a business planning system

and associated workshop designed to embed AI use within community contexts and support the collective capacities entrepreneurs need to sustain their ventures amid an unstable technological landscape.

2.3 Community-Based Participatory Design in AI and Entrepreneurship

Scholars in HCI have long emphasized that design processes risk reproducing existing social inequities if they fail to center the perspectives of those most affected by technological systems [22, 36]. Community-Based Participatory Research (CBPR) [44]—a methodological approach aligned with participatory design [64] (PD)—seeks to equitably involve community members and researchers across all stages of a project. Within HCI, CBPR has been increasingly taken up to ensure that communities most affected by technology play a role in shaping the design, development, and deployment, and through it, foster an equitable design process (e.g., [22, 29, 39, 50, 56]).

Within entrepreneurial contexts, CBPR’s community-centric orientation aligns closely with HCI research that shows how trust, peer network, non-technological efforts, and shared visions are essential in supporting and sustaining entrepreneurs and small businesses [23, 39, 49]. Community spaces such as makerspaces and entrepreneurial hubs [47] scaffold digital engagement [40, 50] and enable practices like help-seeking and peer learning [51, 52] and adaptation to technological change [48]. HCI scholarship draws attention to the value of community-based research in understanding the resources available to the community, the power dynamics that are at play, and the entangled complexities where designed technologies are situated [20, 30, 95].

Aligned with this, scholarship describes a growing “participatory turn” in AI design, calling for approaches that embed communities most affected by AI systems into the processes of design, development, and governance in ways that extend beyond token consultation toward meaningful influence and shared agency [21]. However, this turn is not without its limitations and critiques. Scholars warn that participation in AI design with clarity on why AI is used, and what power participants have in shaping it, can risk cooption and cause harm [14, 27]. Similarly, other scholars highlight the need to build reciprocity, lest the participation become a one-off “blue sky ideation” that privileges unconstrained novelty over the on-the-ground reality that those most affected by technology face [36, 71, 74]. These challenges are further amplified by the vast design space and speculative promise associated with emerging AI technologies [21, 62].

In this work, we take up CBPR in the design and deployment of an AI-embedded business-planning system within a feminist makerspace in Pittsburgh, Pennsylvania. Our approach builds on a prior four-part workshop series on generative AI and entrepreneurship at another local small business hub [48]. Guided by lessons from these workshops and CBPR’s call for building reciprocity, we scoped our design intentionally on a constrained, practical application—business planning—that offers direct, tangible value for local entrepreneurs. Moreover, we rely on our existing relationship with the makerspace to deploy the technology and examine how its adoption can be scaffolded through the community setting. Next,

we describe our system design informed by our prior workshop and grounded in entrepreneurs' needs, the deployment context, and our broader commitment to collective resilience.

3 BizCHAT

3.1 Business Planning as Context for Resilience Building

We chose business planning as the focus for our system because it provides a practical context for resilience building. Business plans are essential, yet often overlooked documents, required for entrepreneurs to apply to grants, loans, and small business accelerators. Beyond accessing capital, business plans are tools for reflection, requiring entrepreneurs to examine goals, strategies, and operations against their expertise. Contrasting with traditional business planning approaches (e.g., SCORE [4]), which emphasize static templates, we take an iterative learning-oriented approach to business planning. Last, we saw that current AI-business planning tools (e.g., LivePlan [42], Upmetrics [1], IdeaBuddy [5]) prioritize automated drafting, rather than accessible interfaces and community connection required for resilience building in practice. Building on these foundations of what makes a business plan, we designed BizCHAT to instantiate community-centered design principles in the context of AI-supported business planning.

3.2 Formative Study and Design Goals

3.2.1 Formative Study. The design of BizCHAT builds directly on a four-part community workshop series and interview with local entrepreneurs focused on integrating (or supporting the refusal to integrate) generative AI into small business workflows [48]. In this workshop series, which focused on practical applications of AI in small business contexts, authors surfaced the long tail of accessory skills (e.g., browser literacy, password management, file literacy, typing skills, clipboard management) that entrepreneurs must acquire in order to effectively use generative AI. In this section, we detail the findings of this formative study and related work from HCI and learning sciences that led to our three design goals for BizCHAT: (1) support diverse levels of digital and business expertise through low-floors and high-ceilings, (2) embed just-in-time learning opportunities that connect AI use directly to business outcomes, and (3) contextualize AI introduction within entrepreneurs' existing knowledge and goals to reduce anxiety and build confidence.

3.2.2 Low Floors, High Ceilings. Previous workshops revealed that entrepreneurs came with highly diverse technical and business backgrounds. Some struggled with accessory skills such as typing or file management, while others were already experimenting simultaneously with many AI tools, such as ChatGPT, DALL-E, and Midjourney, among others. Thus, BizCHAT's first design consideration is to build "low-floors and high-ceilings"—a pedagogical term which refers to designing interventions that are accessible to beginners while providing opportunities for advanced learners to engage in deeper exploration [70]. For small business owners, this means creating a system accessible to users with limited digital skills, while remaining extensible for those already familiar with tools like ChatGPT. As small business owners have diverse levels of typing and writing proficiency—from two-finger typing to touch

typing—we scaffold the editing process so that **Voice-to-Text (2c)** is available in every step: describing the business, dictating edits, and chatting with BizCHAT. For users comfortable with typing, BizCHAT's **Rich-text Editor (2d)** enables more control through direct manipulation of their business plan. As prompting remains a barrier for non-expert users of LLMs [96], BizCHAT creates two **Prompt Suggestions (2a)** for every conversation turn—one prompt suggestion focuses on the current topic (exploitation); the other focuses on a new topic (exploration) [75]. When BizCHAT provides suggestions, the user can one-click **Apply (2b)** suggestions from the chat to the business plan in the rich-text editor. For more advanced users, BizCHAT enables in-line text generation (**2e**), where users can specify criteria to generate text and view exemplars. Further, the **Rich-text Editor (2d)** allows users without word-processing skills to specify text styles and **Export (2f)** to a standard template, without the additional barriers of complex document formatting.

3.2.3 Just-in-Time Learning. Prior design workshops revealed that for entrepreneurs from resource-constrained backgrounds, the value of general-purpose AI is often unclear [48]. Further, participants emphasized that, given their busy schedules, they could not step away from their businesses for AI literacy training or generic AI tutorials. Thus, BizCHAT's second design consideration is just-in-time learning [38, 82]—in situ micro-learning opportunities that support entrepreneurial education. BizCHAT aims to develop skills in three areas: understanding effective business plans, building reflection and help-seeking skills, and effectively using generative AI. To do so, BizCHAT's onboarding flow includes an **Informational Video** reviewing key motivations for business plan creation and detailing their key components. BizCHAT then scaffolds users to iterate on each section of their business plan, suggesting relevant **Examples (2e)** from U.S. Small Business Administration². In addition, BizCHAT provides users with relevant suggested questions as tooltips in the editor for the **Business Plan Assistant (2e)**. Last, to build help-seeking skills and facilitate self-reflection, BizCHAT's **Prepare to Pitch (2g)** provides users a list of questions to ask an expert tailored to their business plan and goal, which are a helpful preparatory step when approaching business coaches for critical feedback.

3.2.4 Contextualized Introduction to Technology. Prior design workshops also revealed how entrepreneurs express anxiety about falling behind or "missing out" on new technological trends, which compounds hesitation to engage with unfamiliar tools [48]. To mitigate such anxieties, introducing technology in the context of users' existing knowledge and goals—rather than focusing on the novel technology—can be a helpful strategy [10, 30]. Therefore, BizCHAT's third design consideration is to contextualize users' interactions with generative AI within their area of expertise—their business [37]. For instance, at each conversation turn, the **Business Plan Assistant (2b)** suggests changes to the business plan in accordance with the user's **Business Plan Goals**, which are explicitly set during onboarding. Further, unlike AI systems that present outputs as a completed document, BizCHAT positions business plans as an evolving document used as a tool for planning and growth. To this

²<https://www.sba.gov/business-guide/plan-your-business/write-your-business-plan>




Figure 2: BizChat supports small business owners to draft and refine business plans. Interface legend. (a) Prompt suggestions for the next turn. (b) One-click *Apply* to insert suggested changes into the plan. (c) Voice-to-text for dictation. (d) Rich-text editor for direct manipulation. (e) Inline generation/exemplars to draft or revise selected text. (f) *Export* to a formatted document. (g) *Prepare to Pitch*: expert-style questions and an option to connect with an expert.

end, BizCHAT seeks to facilitate deeper engagement with a broader community of support (e.g., expert business coaches, an essential part of an entrepreneurial ecosystem [39]). BizCHAT offers users to **Connect with an Expert** (2g) and aims to decrease reputational risks by providing entrepreneurs a list of **Questions to Ask an Expert** (2g) about their business plan. In doing so, BizCHAT acknowledges that, even with the latest technology, social support among minority entrepreneurs is essential [40, 49–51].

3.3 BizCHAT User Scenario

To illustrate how BizCHAT can be used to draft, iterate, and take next steps with a business plan, let us follow Tina, an aspiring entrepreneur with a boutique clothing company. Although she runs her business from her mobile phone, primarily selling through Facebook Marketplace and Etsy, like many other solo-preneurs [84], she finds the maintenance of this constellation of digital tools is tedious and takes her away from core business activities. As a result, keeping pace with rapidly advancing technologies, like ChatGPT or Gemini, is challenging and overwhelming [48]. Moreover, Tina finds it unclear how to extract value from these tools long-term.

3.3.1 The need for a business plan. After three years of consecutive four-digit revenue proving the viability of her small business, Tina wants to expand her business into a brick-and-mortar store in her city’s downtown. While attending a community pop-up market for local artists, her solo-preneur friend informs her about a recent funding opportunity to revitalize her city’s downtown storefronts. However, to apply for the funding and subsidized shopfront that

she needs to scale up her business, Tina, like many other small business owners on the cusp of growth, needs a business plan. While reviewing the requirements for the grant application, Tina finds BizCHAT on the resource page.

3.3.2 Drafting a Business Plan with BizCHAT. After logging in, Tina is directed to BizCHAT’s onboarding, which asks her to enter her website, but also gives her the option to proceed by just chatting if she does not have one prepared. Tina enters the URL of the website her niece vibe-coded for her. BizCHAT then extracts relevant information from the website. But, since the website is incomplete without essential details for a business plan (e.g., product information, legal structure), BizCHAT detects the missing details and asks Tina for the information it is missing. BizCHAT asks: “What’s your vision statement?” Not lacking a vision, but unclear on how to precisely articulate it, Tina clicks “Not sure? Click for some ideas.” to view BizCHAT’s response suggestions (Figure 3.1). After BizCHAT has all of the information needed for a first draft, it generates a first draft of Tina’s business plan. While waiting, Tina watches a just-in-time learning module about the sections of a business plan. While reviewing the business plan BizCHAT generates, she notices an error in the Executive Summary. Because her two-finger typing is less efficient, she uses BizCHAT’s voice-to-text to edit the plan in the rich-text editor. BizCHAT suggests changes, and after reviewing, Tina clicks “Apply” and watches BizCHAT make the changes to her business plan, highlighting exactly where it is making the change. BizCHAT’s prompt suggestions (Figure 2a), then suggest making some changes to the Market Analysis section. To verify the changes




Figure 3: BizChat scaffolds the onboarding process for business plan creation. To (1) collect business context, users can enter their website or chat with the assistant about their idea. Next, users (2) specify the purpose of their business plan, such as starting a new venture, applying for funding, or expanding an existing business. Users can optionally (3) watch an informational video that introduces business plans and the BizChat workflow. Finally, BizCHAT automatically generates a first draft of the business plan based on the provided information.

suggested, Tina compares with an example Market Analysis small business plan from SBA.gov [7] in the BizCHAT interface.

With Tina’s newfound understanding of what an effective market analysis section looks like, she presses “Apply” and applies the changes. While Tina has no interest in complex file-formatting for business plans, she clicks “Export” to get a pre-formatted business plan in a standard format to share with others. With a draft business plan in hand, Tina wants to prepare for the grant application. She navigates to the “Prepare to Pitch” tab in the BizCHAT interface. Previously apprehensive to share her plan with her mentors, she reviews BizCHAT’s custom “Questions to Ask an Expert,” and feels more prepared to share her business plan before applying for grants.

3.3.3 Overcoming technology gaps with her local network. With a draft in hand, Tina shared her business plan with her close entrepreneurial friends. While she felt clear about her product and vision, she was unsure how to build the financial projections required for the grant. Tina’s mentor walked her through revenue and expense estimates, filling in gaps Tina had struggled to overcome alone with BizCHAT. Meeting the requirements for the grant application, Tina is able to apply to the revitalization grant and explore more opportunities in BizCHAT’s **Explore Local Grants** tab (See Figure 5). Encouraged by this experience, Tina shares BizCHAT with other trusted entrepreneurs in her network.

3.4 Implementation

BizCHAT is implemented as a React application built on Next.js, with Firebase handling authentication, data storage, and telemetry. BizCHAT is built on top of TipTap for its rich text editor and is deployed with Vercel. At the core of BizCHAT’s business plan generation is a structured business plan JSON that is progressively filled out during onboarding. Each field is initialized with a NOT_FOUND placeholder and is dynamically updated as the user provides information. The onboarding assistant, powered by GPT-4 Turbo

Preview, asks only about missing fields and adapts questions accordingly. At each conversation turn, an observer LLM powered by GPT-4 extracts business information from users’ answers to onboarding questions and updates the business plan JSON.

When a website is provided, BizCHAT crawls within-domain pages and uses GPT-4 with function-calling to extract structured information. The onboarding assistant dynamically adjusts what questions to ask based on what information is missing from the business plan JSON. If no website is provided, the chat begins immediately. To ensure enough context is provided to each business plan prompt section, minimal completion requires a name, mission, target market, at least one product with a value proposition, and a basic marketing and sales strategy.

Once the structured business plan JSON meets completion criteria, BizCHAT makes asynchronous calls to gpt-4-turbo to generate each section, with only relevant structured information inserted into the prompt. This approach maintains interactive latencies (around 8–15 seconds per section) instead of waiting for a single long response. Each business plan section’s prompt is few-shot with examples provided from SBA.gov exemplar business [7]. BizCHAT’s business plan assistant is powered by **GPT-4o** and is prompted to generate <suggestion> tags, which render in-chat components that allow users to directly apply suggestions to their business plan by clicking a button. We intentionally did not implement automated evaluation metrics. Instead, to reinforce accountability with our community partner, the generated plans were shared with community partners for a sanity check, focusing on practical utility in local contexts.

4 Methods

4.1 Community Site

Our research was conducted in partnership with Prototype PGH (or “Prototype” for short), a feminist makerspace located in Pittsburgh, Pennsylvania. Founded in 2016, Prototype’s mission is to foster

gender and racial equity in creative entrepreneurship by “providing affordable access to high tech tools and equipment, offering workshops that prioritize the experiences of marginalized communities, and cultivating a professional support network.” Prototype’s ethos is that “everything is a prototype,” emphasizing feedback and iteration as essential to early-stage creative and entrepreneurial practice.

Prototype functions as both a makerspace and an entrepreneurial hub, providing shared equipment (e.g., laser cutters, 3D printers, and sewing machines), skill-building workshops, and a supportive peer network. Its programming is intentionally designed to serve underrepresented entrepreneurs, particularly women and gender minorities, and people of color, who often face compounded barriers to accessing resources, technical skills, and professional support.

Our collaboration with Prototype was built on longstanding relationships between the members of the research team and the community site. One author has been a dues-paying member and collaborator for several years (2017–2025), participating in both the makerspace and its business incubator programs. This positionality (see subsection 4.8) facilitated trust with Prototype leadership and participants, and allowed us to ground the design and deployment of BizCHAT within a community already invested in inclusive, iterative approaches to entrepreneurship.

Participants were recruited through the authors’ existing local networks, developed over years of community-based research in entrepreneurship and technology. Recruitment materials were circulated via personal and professional networks on platforms including LinkedIn, X (formerly Twitter), and via Prototype’s existing communications channels (i.e., Instagram, newsletter), extending reach to its broader membership and entrepreneurial community.

4.2 Four Workshops (N=21)

We developed and deployed BizCHAT through a series of four workshops hosted in collaboration with Prototype between January and August 2025. The first workshop (January 2025) was integrated into Prototype’s annual incubator program as part of a session on financial planning. The two-hour workshop was structured in two parts: during the first hour, facilitators introduced key financial documents and strategies for small business owners. One author was present throughout onboarding to provide live support, helping participants navigate technical issues and ensuring they could begin drafting their business plans with the tool.

This framing was a deliberate methodological choice. Rather than centering AI in our outreach and workshop titles, we positioned business planning as the focal point. We did so to create an accessible entry point for entrepreneurs who might feel anxious about AI-focused framing, an issue highlighted in prior work [48], to ground engagement in participants’ existing business expertise, and to observe how entrepreneurs engage with AI on their own terms within a task they already needed to complete. We reflect on the tensions this choice introduced—including concerns about transparency raised by one community partner—in Section 6.2.

Following positive reception, we were invited to return for three additional standalone workshops dedicated to BizCHAT and business planning. We hosted three workshops that took place on Saturdays in June, July, and August 2025. Each of these sessions followed

a consistent structure under the title “Write Your Business Plan: Resilience and Uncertainty Planning Workshop.” The workshops began with content on the basic structure of a business plan and on the role of business planning as a resilience practice, positioning business plans as living, evolving documents that can adjust based on the needs of the business. We explicitly framed business plans as drafts requiring ongoing iteration, emphasizing that participants could create different versions for different goals (e.g., grant applications vs. loan proposals). This framing was intentional. By positioning the business plan, and not the AI technology, as the workshop’s focal point, we sought to ground participants’ engagement in their existing business expertise rather than foregrounding technical concerns about AI.

After introductory content, participants were invited to use BizCHAT to create a draft business plan. As in the first workshop, a research team member was present to answer questions and help workshop participants overcome technical difficulties while using the tool. Rather than beginning with abstract explanations of LLM limitations (e.g., hallucinations, training data biases), we relied on participants’ expertise to surface issues as they arose during use. For instance, when participants noticed AI-generated content that did not align with their business knowledge—such as incorrect product descriptions or generic market analysis that did not reflect their actual customer base—these observations prompted group-wide discussions about AI accuracy, the importance of human oversight, and strategies for evaluating output quality. This organic approach to discussing AI limitations allowed participants to identify issues relevant to their specific contexts rather than requiring them to navigate abstract technical terminology before understanding its practical implications.

Following their use of BizCHAT, participants engaged in reflection activities with paper-based planning tools that scaffolded resilience strategies such as if–then planning (shown to strengthen adaptive responses under uncertainty [32]), goal setting, reflecting on challenges, and writing down next steps for accountability. These post-use activities were designed to complement BizCHAT by encouraging participants to identify concrete next steps beyond the AI-generated draft (e.g., connecting with mentors, researching funding opportunities, refining financial projections with expert support).

4.3 Follow-up interviews (N=10)

Participants who attended a workshop were invited to opt in to a follow-up conversation. We conducted semi-structured follow-up interviews with eight workshop participants and two entrepreneurial support personnel in the weeks following each session. Interviews served two purposes: (1) to understand what steps participants took after the workshop based on the goals they articulated during in-session reflection activities; and (2) to offer an accountability touchpoint and additional support. Interviews lasted 30–60 minutes, were scheduled at participants’ convenience, and were conducted remotely via Zoom. We started interviews by asking questions to situate participants’ business context within the workshop (e.g., “What are your biggest business challenges right now?”, “What inspired you to attend the business planning workshop”) and moved

PID	Business Description
P1	Designer and creative director; freelance design business owner
P2	Co-founder, makerspace for neurodivergent vocational training
P3	Co-founder, makerspace for neurodivergent vocational training
P4	Gift baskets and party planning business
P5	Custom jewelry and apparel designer
P6	Psychic and medium
P7	Nuisance control business; youth mentor at Job Corps
P8	Founder, consulting firm advising micro-businesses and nonprofits
P9	Director of a local makerspace; writer and novelist
P10	Licensed counselor & missionary

Table 1: Interview participants’ businesses and roles. Businesses span creative industries, community organizations, service sectors, and individual practices.

towards questions to probe resilient behaviors exhibited by entrepreneurs (e.g., “Have you asked anyone for feedback on your business plan?”). With consent, interviews were audio-recorded and transcribed for analysis. The protocol prompted participants to revisit their written goals, describe progress and barriers, reflect on any subsequent use of BizCHAT or other tools, discuss peer supports they engaged, and share ethical concerns or refusals related to AI use. The protocol prompted participants to describe progress and barriers towards their business’ goals, reflect on any subsequent use of BizCHAT or other tools, and consider how the workshop helped or hindered their next steps. We also offered resource referrals and optional check-ins where appropriate to sustain momentum.

4.4 Informed Consent and Informed Refusal

We separated tool use from research participation so that entrepreneurs could access BizCHAT without opting into the research study. At workshop onboarding and on the BizCHAT site, participants encountered a consent modal (see Figure 4) with three clear choices: “Opt in,” “Opt out,” and “Learn more.” Consistent with work on informed refusal [12], “Opt out” allowed full use of the tool without contributing data to the study; “Opt in” permitted collection of de-identified interaction data and artifacts for research; and “Learn more” expanded details about the study, the research team, and the data collected [85].

To mitigate power asymmetries that can fuel tensions in community–academic partnerships, we disclosed team identities and affiliations, what data (if any) would be collected under each choice, who would have access, storage location and duration, and points of contact for questions or withdrawal [85]. Participants were reminded that non-participation would not affect access to the tool or workshop activities, and they could change their choice at any time. If participants revoked consent, we were prepared to delete previously collected data from research storage to the extent permitted by law and platform constraints (although there was no occurrence



Figure 4: We intentionally separate tool access from research participation to enable informed consent and refusal, mitigate community–academic power asymmetries, and ensure entrepreneurs can fully use BizCHAT regardless of research choice; the modal offers “Opt in,” “Opt out,” and “Learn more,” with links to plain-language details on purpose, data handling, and withdrawal.

of this). Of the 48 users on the BizCHAT platform throughout the duration of the study—recruited exclusively through word-of-mouth and workshop advertising—18 chose to opt out.

4.5 Participants

Participants in our study included small business owners and entrepreneurs recruited through local networks and community channels. While recruitment was supported by the feminist makerspace, participants were not necessarily members of the space. The group reflected the broader entrepreneurial community connected to it. To respect privacy and avoid re-identification within a relatively small community [6], we did not collect nor will report granular demographic information. However, participant demographics broadly reflected the inclusive orientation of the makerspace, centering women, non-binary people, and people of color in entrepreneurial practice. More details for interview participants in particular can be found in Table 1.

4.6 Data Analysis and Evaluation

We followed a critical incident technique to identify specific interactions or events that represented significant moments of success or challenge during the participants’ interaction with BizCHAT [26]. To this end, the first author wrote analytic memos immediately after each interview and workshop. The notes captured conversation details, impressions, and critical incidents—such as moments of confusion, breakthrough, or successful application of suggestions—directly informing our research questions: revealing barriers and supports to adoption, illustrating how community infrastructures mediated use and refusal of AI tools, and surfacing behavioral




Figure 5: BizCHAT’s *Explore Local Grants* tab provides a curated list of local grants and loans maintained by the research team. Users can search, filter by deadlines, and view detailed information about each opportunity, including funding amount, equity requirements, and application dates. This feature helps entrepreneurs discover and apply for relevant opportunities without leaving the platform.

and cognitive strategies entrepreneurs employed to build resilience. All authors reviewed field notes, transcripts, and memos to identify and code these critical incidents. These codes captured different types of challenges and successes, such as moments of deep sensemaking, quick drafting of a plan, critical concerns with AI, or moments of frustration with the tool. Through iterative comparison and discussion among the research team, we converged on themes related to research questions (e.g., barriers to accessing capital, tension between sensemaking versus instant results, and the importance of communal scaffolds to support AI literacy and planning).

After synthesizing themes, we conducted a member check with our community partner at Prototype PGH. In this session, we presented a concise set of emerging findings with illustrative quotes and invited critique, corrections, and elaboration. Following guidance from community-engaged HCI practice [85], we treated member checking as both to establish the validity of our findings and an accountability practice to further our ongoing partnership, rather than a one-off evaluation. The refined codes form the basis of the findings presented below.

4.7 BizChat’s In-App Data Collection

4.7.1 Pre and Post-Use Surveys. For consenting BizChat users, within the interface, we administered a pre- and post-survey at two critical moments: 1) before participants generated their business plan and 2) before participants exported their business plan. This pre-survey captured baseline measures of participants’ entrepreneurial confidence and comfort, and uses cases with AI tools.

The post-survey repeated these measures and added an option for entrepreneurs to receive expert feedback from a community member on their business plan. We report the survey and averages from responses in the Appendix.

Completion of either survey was not required to use BizCHAT and participants could skip it with a single button and continue directly to tool use or export. This design choice was intentional for opted-in participants to further mitigate power dynamics common in community-based academic partnerships, although it did lead to low survey completion rates. We tracked completion rates for the pre- (10 responses) and post- (4 responses) survey, and used them to complement the qualitative data analysis.

4.7.2 Log Data. We analyzed application logs for the workshop deployment period (January through August 2025). We implemented our telemetry schema via the Firebase SDK, and it evolved over the deployment; when metrics required fields were introduced mid-study, we report them on the eligible subset of users for whom those fields existed and explicitly indicate this denominator (denoted by “eligible” in Findings). Consistent with our consent design, opted-out users were excluded from analysis. Primary metrics include number of users (total vs opted-in), sessions and capped session length, survey completion rates (pre among opted-in; post among pre completers), editor applies and inline-edit sessions, and AI prompt suggestion-acceptances. For our analysis, we explicitly filtered out non-participating users, test accounts, and accounts with no business plan. To calculate session duration, we computed capped session durations by summing consecutive event gaps within each session, capping each gap at five minutes and then averaging across sessions.

4.8 Positionality

We situate ourselves within the community, region, and research traditions that shape this work. We do so to acknowledge how our identities, histories, and relationships informed design choices, data collection, analysis, and claims.

The first author is an entrepreneur who grew up in the region—he is a co-founder of Y-Combinator-backed technology-based, direct-to-consumer business. His proximity to participants geographically and familiarity with the region support rapport and trust. He served as a facilitator during the first workshop and led the analysis. The second author is interested in issues of power and inequity and has been conducting research in the region, engaging tangentially with the makerspace through a separate workforce development project. Similarly to the first author, he had created a tech startup in 2014. He approaches AI with skepticism; his involvement here stems from the last author’s introduction to its potential applications in this space, though he continues to approach the technology with caution and curiosity. The last author sought out the feminist makerspace during her PhD to find a like-minded and critical community with diverse membership and less-narrow conceptualizations of “hacking.” She has been connected to the community space by being first a dues-paying member, and then a formal collaborator for the last eight years. In addition, she has worked closely with over 200 entrepreneurs through her academic and industry work experience, primarily focused on the needs and aspirations of underrepresented entrepreneurs. Outside of research, her creative

pursuits include professional dance, which further connects her to the makerspace's diverse forms of making and expression.

Our backgrounds enabled rapport and helped position us as relative insiders who understand the pressures of business planning and capital access firsthand. However, our academic positions mean we do not share the financial precarity that shapes necessity-driven entrepreneurship for many participants.

5 Findings

In this section, we detail how BizCHAT helped entrepreneurs translate ideas into the “language of business,” quickly generate first drafts, and edit plans in their own voice. Yet, this speed sometimes bypassed deeper business sensemaking, surfacing a tension between efficiency and reflection. In addition, we share how participants turned to peer networks to refine their plans, build AI literacies, and seed their own community infrastructures to negotiate adoption.

5.1 Leveling Barriers to Accessing Capital and the “Language of Business”

Overall, participants viewed BizCHAT as a tool that levels the playing field for accessing capital by: (1) translating entrepreneurs' ideas into the “language of business” and (2) shortening the time it takes to create a first draft.

Participants valued that BizCHAT, “translated everything to the right speak” (P1). P2 and P3, who are starting a makerspace for vocational training for neuro-divergent individuals, shared how BizCHAT reflected their words back in recognizable terms: “and then to have it regurgitated back to me in like business lingo, which is not my expertise was like, thank you. ...that's why I really enjoyed it.” P2 framed this translation as navigating a gate-keeping language that previously felt inaccessible: “*We are not business people, but that levels the field. So if someone has an idea like, I've got this great idea, but I'm not a business person, if you have the AI generators ...it's a hurdle that somebody doesn't have to take, you know?*” For P2, “hurdles” meant taking unstructured thoughts from “*writing whiteboards*” and “*typing pages*” into polished formats. This was reflected in usage logs: across 13 eligible users' onboarding conversations, participants engaged in an average of 6.54 user turns (median 7, max 12), and crucially, they clicked “*Not sure? Click for some ideas*” (see fig. 3.1) prompt-suggestions 3.27 times per active user—nearly half the median onboarding chat-turns. This shows that during onboarding, participants leaned heavily on AI-generated suggestions to surface ideas and translate them into business language, a concrete instantiation of BizCHAT's low-floor, high-ceiling design that helped them articulate what they meant even when they were unsure.

For P4, BizCHAT lowered the floor by quickly consolidating information for her. P4 shared how gathering previously scattered business information from her website (e.g., product line, mission statement copy) reduced the effort needed to assemble her plan and lessened the stress of tracking down and recalling those details on her own. P4, who self-describes as a “*grant-queen*”, but “*not a techie*” now felt empowered to take the information from BizCHAT, and use it for other grant applications: “*The information that I learned*

from the BizCHAT, I now can take pieces of that... and then just pop it in, instead of looking through all my papers.”

Participants also valued the agency BizCHAT, unlike generic text interfaces with un-editable outputs, afforded them to reshape that “*right speak*.” As P4 put it, “*It allowed me to see it, and in sections, and then I can work with it that way ...I needed to reword it in a way that sounds like I wrote it, and not a robot wrote it*” (P4). But, building the agency to edit the business plan was not purely due to interface design, but also in how the research team and community partner positioned business plans as living, evolving documents during the workshop series. For P2, this was a major un-blocker, “*I didn't know the plan could be flexible ...I thought once you had a plan, it was carved in stone.*” This framing in the workshop was a helpful affirmation that P2 could adapt her plan as her business adapted.

In addition to translating entrepreneurs' ideas into the language of business, participants also expressed how BizCHAT significantly reduced their time to get a first draft. Both P4 and P5 reported the time it took to draft a business plan went from “*six months*” to a matter of “*an hour*” (P4) or “*minutes*” (P5). As P5 put it, “*It took me six months to write my last business plan... [to make sure] all the I's were dotted, all the T's were crossed.*” For P5, who runs a custom jewelry and apparel business, building this momentum to get a first draft had a very tangible outcome. Later, she reported using the business plan BizCHAT helped her draft to apply, and ultimately get accepted to a major international fashion event. Even if the draft was imperfect, BizCHAT provided the momentum she needed—something tangible to share with a friend—which ultimately set her on the path toward achieving a major milestone for her business.

For P6, who self-identified with ADHD, speeding up the drafting process intersected with equity. Unlike domain-agnostic tools, BizCHAT provided a non-blank starting point: “*I have ADHD... it was super helpful... it's hard to get from zero to something going*” (P6). She described the tool as moving her “*to a starting point that's like further down*” and as “*add[ing] like the equity element*” toward parity with those without ADHD (P6).

For P7, reducing the time to create a first draft helped reframe the emotional stakes of sharing his business plan. Previously, P7 viewed sharing his business plan as exposing himself to the risk of IP theft from possible lenders. For instance, P7 recalled this anxiety around sharing plans: “*If you turn me down, you have my plan... you can just give it to your kid, and they can go off and be rich... I just handed it to you on a piece of paper. Here you go! Here's all my hard work*.” P7, reflecting on his past experiences as a member of a marginalized community, described how sharing his business plan once carried a form of extraction anxiety—the fear that lenders could reject him yet still profit from his hard work in creating the plan. When business planning required months of painstaking effort, sharing that plan felt like handing over deeply personal labor that could be exploited. By reducing the time and emotional investment required to produce a draft, BizChat shifted what felt vulnerable: the business idea itself remained valuable, but the document representing it no longer carried the weight of months of personal struggle. Importantly, this shift reflects a change in participants' *subjective assessment of risk*, rather than any change in the objective legal ownership or factual risk of IP theft. As P7 explained, when a draft takes an hour rather than six months to create, the anxiety of potential theft diminishes—not because the IP is less valuable, but because the personal cost of

producing the artifact to share is dramatically lower. Moreover, the ease of iteration meant P7 could quickly generate updated versions or tailor plans for different audiences, reducing the sense that any single document represented an irreplaceable, static record of his ideas that could be stolen and exploited.

BizCHAT helped entrepreneurs overcome barriers by translating their ideas into the “*language of business*” and accelerated their drafting process, while still supporting agency in shaping their plans. Yet, while participants valued these efficiencies, they also grappled with the balance of quickly drafting a plan, and engaging more deeply in making sense of their business. Thus, in the next section we highlight this tension.

5.2 Business Sensemaking vs. Instant Results

Entrepreneurs (P4, P5) were blunt that time is scarce and business plans are boring. The pull to get an instant draft and move on was real and reasonable, especially given time and resource constraints. However, in the context of our resilience planning workshops and deployment of BizCHAT, this surfaced a tension between the need for instant results, while supporting long-term planning and business sensemaking.

As mentioned in the previous section, quick results were celebrated because they removed the drudgery of business planning. As P4 put it, “*I don’t have time to sit down and write out a business plan... it’s so freaking boring to create step by step.*” For P4, who had previously spent six months creating a business plan, the instant results of the BizCHAT workshop, provided the necessary momentum to get over the boredom of business plan writing.

However, entrepreneurial support personnel were keen to share how quick drafting can mask gaps in preparedness. As P8 put it, “*If [BizCHAT] makes a plan that helps people get money, [but] they have no business owning a business and they don’t have the skills, [the plan] doesn’t help.*” (P8) She also cautioned against over-trusting tidy output: “*One of the issues, where [BizCHAT] can’t find the information on the website: Is it making something up? It sounds compelling, so people are not going to challenge it.*”

Through low-floor high-ceiling design principles, BizCHAT started, but failed at fully navigating this tension. Through onboarding, entrepreneurs are asked to think through many aspects of their business that ultimately appear in their business plan. P8 highlighted how this can be powerful for self-reflection, “*If they survive the gauntlet of being asked all of the questions, that could be really helpful to people.*” P8 expanded to say that, in being asked these questions, the entrepreneur could have realizations about themselves—what if they realize they actually do not want to start a business?

Participants also described onboarding as just-in-time learning that turns unknown unknowns into known questions. Contrasting with general-purpose AI tools like ChatGPT, P6 noted the importance of BizCHAT’s guided flow over open-ended prompting: “*with ChatGPT, it feels like I’m guiding the tool, whereas like BizCHAT feels like it’s guiding me.*” P6 added: “*ChatGPT...it takes more brain power to even know what to ask it.*” Building on that, P2 shared that “*It’s telling me what I don’t know, when I don’t know.*” P3 emphasized how the tool surfaced blind spots, describing it as something that “*takes everything I’ve learned and polishes it and clarifies it and make sure*

I think about things that maybe I hadn’t thought about.” She linked this reflection to a broader goal of staying competitive, noting that “*the tools can help you learn about what the questions are, so that you stay competitive and that you stay current.*”

Beyond surfacing unknown unknowns, for P6, the scaffolding was a helpful tool to shift away from “shyness” and get her in the mindset of running her business: “*I’m someone who shies away from dreaming bigger...It helped me get in the mindset of running my business.*” For P6, the structured prompts asking her to articulate her target market, value proposition, and growth strategies forced her to inhabit the role of a business owner—thinking strategically rather than operationally. By centering her expertise and asking her to make decisions about her business throughout the onboarding process, BizChat positioned her as the authority, which shifted her self-perception from someone “dreaming” to someone actively planning and executing. P9, who once went through the Prototype PGH Incubator program and now runs communications at a local makerspace, observed this pattern across many entrepreneurs, questioning whether the biggest barrier for entrepreneurs is less about technical knowledge and more about developing the entrepreneurial mindset: “*Do people need more training in having an entrepreneurial spirit?*” P9’s question highlights a critical tension: while BizChat’s scaffolding helped P6 adopt a strategic mindset by prompting her to make concrete business decisions, it remained unclear whether this was genuine mindset development or merely performing the role of entrepreneur through AI-prompted responses. P9’s skepticism—echoed in her earlier concern about whether AI-generated plans truly “train” entrepreneurs—suggests that quickly producing a polished plan might substitute for, rather than cultivate, the deeper entrepreneurial thinking that comes from wrestling with business challenges themselves. In this way, BizCHAT’s scaffold encouraged reflection, learning, and goal setting, yet the question of whether AI-assisted planning builds or bypasses entrepreneurial capacity remained unresolved.

For entrepreneurs without a website, BizCHAT asks questions about each section of their business plan before creating a draft. But if an entrepreneur’s website already contains enough relevant information, BizCHAT skips these questions and generates a draft immediately. In this case, the entrepreneur may not encounter the “*gauntlet of questions*” that P8 describes. P1’s case illustrates this miss. He arrived with a website, so BizCHAT produced a complete plan immediately. What drew him to the workshop, however, was the promise of resilience planning as an ongoing practice, not just a document. He warned that over-ease can reduce ownership: “*if you make things too easy, people don’t value it enough... if you get to the finish line, you have less ownership over it.*” P1 had recently transitioned back to running a freelance design agency after a round of layoffs at his previous employer. For an operator with fifteen years in business who wanted to think through the right contingencies, BizCHAT’s low-floor and high-ceiling approach did not sufficiently slow down or deepen reflection. Log data tells a similar story. After onboarding, users did not significantly iterate on their business plan, with users (20 eligible) averaging 0.77 rich-text edits (see Figure 2d) and 0.77 AI-applies per user (see Figure 2b). Yet, despite limited post-onboarding edits, users (26 eligible) averaged a session length of 32 minutes, with a median of seven onboarding chat-turns—suggesting users spent substantial time reflecting and

engaging during the guided onboarding flow, even if they did not edit their business plan after.

While instant results provided necessary momentum and lessened the drudgery of planning, they also risked masking blind spots. The tool helped to uncover unknown unknowns and foster reflection, yet in cases like P1's reveal risks when drafting outpaces deeper engagement. Both within and beyond our workshop series, entrepreneurs turned to peer support to develop AI literacies to navigate this tension, which we detail in the next section.

5.3 Communal Scaffolds for AI Literacy and Resilience Planning

After the workshop, peer support helped to take participants' business plans to the finish line. For P5, continued peer support after the BizCHAT workshop was crucial to turning her plan into a final draft. She recalled how a friend helped stitch together financial projections to finalize her plan for an application, and she ultimately used the BizCHAT-started plan to apply and get accepted to a major international fashion event. Other workshop participants, P6 and another, met up bi-weekly to continue refining their plans together.

While some participants described ad-hoc meetups to take next steps, others wanted to turn this into more formal community infrastructure. For instance, both P7 and P4 suggested hosting their own variations of the BizCHAT workshop within their local communities. P4 expressed interest in hosting her own version of the BizCHAT workshop within her immediate social circle, describing plans to run an informal, small-group session out of her garage to share the tool and process with friends who were curious but hesitant to engage through formal programs. For P4, who emphasized the importance of in-person contact and relational support during the workshop, it mattered that *she* would be the one creating this intimate environment. She envisioned hosting in a way that would foster connection and comfort: *"When I do it, I want to be able to feed the people while we're there. But I want to keep it like to 10 to 15 people, because I think that's a good size."* (P4).

P7, who mentors young adults through a federal workforce training program and runs a nuisance control business, explicitly wanted to host BizCHAT workshops in places like his local public library and local entrepreneurial hub with small cohorts. He described this desire as part of a goal towards supporting greater AI-literacy in his community, likening it to recipe cards: *"I want to introduce literacy with AI—like Betty Crocker sent out recipe cards."* But, here, for P7, *how* AI is introduced mattered. He contrasted BizCHAT's workshops with other academic workshops, in the sense that it disseminated digital literacy at a point focused on practical application rather than new technology: *"What I liked... was 'How can I help someone who obviously needs help right away'... it wasn't like some 'We're gonna help you start a business' type of thing."* For P7, as he plans to bring this to his own community, they emphasized the need to have a trusted messenger. He wanted to host the workshop himself because *"If I did it, they'd believe it."* Here, P7 signals more than personal pride—he emphasized that, in his view, AI literacy is relational. P7 framed information sharing as an act of trust, where credibility depends on *who* cross the boundary to bring back new knowledge. P7 described this in the metaphor of "going to get the bread," drawing on histories of someone risking themselves to bring

resources from inside slave-owners homes back to the field. In the context of AI, they see themselves as the one who must "run in, grab the information, and run back out," testing new tools before introducing them to others.

P7 grounded this in Pittsburgh's geography and culture, reflecting on how a local main street acts as both a physical and symbolic divider. *"Some people will say it's the train tracks. Which, in fairness, it is."* they noted, pointing out how these divides shape whose knowledge is trusted. P7 shared how his tattoos of a Pittsburgh landmark were his proof of credibility—his way of signaling, *"I'm from here, I'm one of you."* Here, his comment that *"if I did it they'd believe it"* underlines that information from *"across the tracks"* may not carry the same weight, as if they were the messenger.

Despite P7 wanting to co-opt the BizCHAT workshop for his own community, our intentional framing of BizCHAT as a business planning workshop, rather than an AI workshop, created tension. For instance, P9 saw us not directly calling out AI in advertising materials, as normalizing use of the technology: *"I do think there's a value in being awake to the ways things are kind of slowly normalized ... when complete abstaining doesn't happen, does it set the stage for normalizing the use of AI in other places?"* P9, who self-describes as *"not a fan of AI"*, and has *"never even opened ChatGPT,"* for *"reasons, political, environmental, [and] creative"*, views this normalization as harmful. Thus, in P9's view, not clearly centering the use of AI in the tool, removed agency from entrepreneurs to engage in refusal of use.

While participants like P4 and P7 articulated intentions to act as trusted messengers introducing BizCHAT within their communities, this study does not observe whether such introductions lead to subsequent community adoption or sustained AI use. Taken together, these accounts show how BizCHAT workshops seeded peer-led infrastructures around business development and AI adoption within the entrepreneurs we interviewed. Yet, as participants considered hosting or adapting the workshops, questions emerged about *how* AI should be introduced, *who* should introduce it, and whether this normalization of use itself is problematic.

6 Discussion

Our findings reveal that resilience emerged from the interaction between BizCHAT, community infrastructure, and peer relationships. We observed this in three interconnected ways. First, entrepreneurs transformed their uncertainty about AI—what P3 noted as "things that maybe I hadn't thought about"—into addressable questions. What began as fear of falling behind became specific, actionable gaps in their business plans. Second, participants built peer-led infrastructures that extended beyond our intervention. P5 needed a friend to complete financial projections. P6 met bi-weekly with another participant to refine plans. P7 and P4 independently planned to host workshops introducing BizChat in their own communities. Third, participants collectively negotiated when to adopt, adapt, or refuse AI outputs, where they developed shared strategies for editing, contextualizing, and evaluating them. Critically, this included legitimizing refusal as informed choice as debates about AI's appropriate role (exemplified by P9's transparency concerns) became sites for collective sensemaking. These observations ground the design implications that we discuss below.

6.1 Designing for Resilience

The forms of resilience above—transforming uncertainty into actionable questions, building peer infrastructure, collectively negotiating AI use—required conditions that our deployment partially created. Reflecting on our experience, we surface three design implications for intentionally supporting such resilience: productive friction to scaffold reflection, communal scaffolds to support peer-based AI literacy development, and co-optability to enable communities to appropriate tools on their own terms.

6.1.1 Productive Friction. Previously, we shared how while BizCHAT aided participants by leveling barriers to accessing capital and the “language of business,” participants often traded their own reflection with instant drafting. In this way, BizCHAT’s low-floor, high-ceiling design principle was successful in lowering barriers to creating business plans, moving from “months to minutes.” However, our log data showed that most entrepreneurs did not iterate significantly on their business plan. This suggests that many participants are settling for “good enough” AI outputs without deeper engagement. This behavior stands adjacent to prior work in HCI, which frames AI-resilient interfaces as those that support users in noticing, judging, and recovering from AI errors [31]. The passive acceptance and the absence of friction or pause for users to question the AI-generated content may risk undermining resilience-building practices, as such reflections are central in creating a strong adaptive capacity [17].

We argue for enabling “productive friction” — moments of intentional pause that prompt users to reflect, revise, or make sense of AI outputs [31, 46]. We draw inspiration from learning sciences literature that advocates for productive struggle and failure in a supportive environment to foster deeper engagement and skill-building [46, 65, 93]. This idea also aligns with the notion of seamless design [43] and design friction [35] within HCI, both of which posit the benefit of sometimes creating friction in the experience to deepen reflection and understanding. Taking these ideas together, we advocate for productive friction, which entails building into the design deliberate difficulty that helps users to understand and improve their judgment of the system’s output.

BizCHAT partially operationalized productive friction through the onboarding flow, where participants were asked a series of detailed questions about their business before receiving AI-generated outputs. This served as a moment of reflection before the AI gave out the results. Reflecting on our experience, we note that deciding where and how to introduce productive friction is a power-laden design choice. It encodes assumptions about what kinds of struggles are valuable. Rather than predefining where users must slow down, designers should offer ways for users to signal whether they seek quick drafting or deeper reflection, and align friction with their goals. Building from work in CBPR, which emphasizes participant agency and cautions against paternalistic design choices (e.g., [8, 28, 90]), we argue that this friction should be *negotiable* as a key consideration is who *wants* to engage in the productive friction and who *needs* quick results. For entrepreneurs facing a large number of tasks and multiple demands on their time, mandatory friction can be exclusionary. We advocate for designing such that users have a choice when to deeply engage and when to prioritize speed.

In BizCHAT, we implemented this negotiability in the onboarding flow (Figure 3) by allowing users to skip it completely. Yet, in

practice, most chose *not* to skip it and went through the detailed onboarding. Some even welcomed the friction. For example, P1, who initially skipped the onboarding flow, later expressed disappointment because that reflection via business planning was what he was seeking. Tied to our low-floor, high-ceiling design principle, we argue that enabling negotiable friction will allow users who value reflection and if their context allows, will engage with the friction, whereas those who need quick output can proceed without being blocked.

While designing for *negotiable productive friction* helps respect user agency, it has the risk of users consistently choosing speed over reflection, and thereby building over-reliance on AI. Building on Glassman et al. [31], we argue that systems (or the context of their deployment) must scaffold noticing, judging, and recovering practices with the output to promote AI literacy: helping users recognize when an output is “good enough,” when it should be challenged, and how to act on either decision. Future work could explore how to design friction points that help users notice and recover from LLM-specific errors, such as hallucinated details or generic outputs that do not reflect their actual business context. As we highlighted in Section 5.3, in practice, BizCHAT users turned to peers to support iterating on AI-generated drafts, and taking next steps with their business plan. Building on this observation, we turn to our next design implication.

6.1.2 Communal Scaffolds. As we touched on above, low-floor, high-ceiling design made it possible for entrepreneurs with diverse digital skills to draft business plans, but it was the communal context that ultimately enabled them to exercise agency, negotiate with AI outputs, and build AI literacy together. In our deployment, lowering the barrier to generating a business plan created opportunities for help-seeking and peer collaboration. For example, after using BizCHAT, participants met with peers over coffee or refined their plans with mentors. It gave them a concrete reason to reach out. BizCHAT also became a vehicle for supporting AI literacy: two participants independently sought to introduce the tool to their own communities, with P4 explicitly framing this as a way to teach others how to use AI for business planning. Reflecting on these experiences, we see the value of carefully deployed AI systems to enable users in developing AI literacy, collective sensemaking, and negotiating the meaning of system outputs.

We argue that AI systems should be designed and deployed to amplify the communal scaffolds already present in the context. In our case, for example, P5 could generate a draft with BizChat but needed friend’s domain expertise to build the financial projections her grant application. Looking at the output, P4 valued BizCHAT’s translation into “business language” but sought peer feedback to make it sound like her, which was an authenticity judgement that benefits from others who know her voice. This is a move to support users to construct an understanding of AI and its outputs together; a way to overcome individual limitations and build collective resilience. Prior research in HCI and entrepreneurship context supports the importance of building such communal scaffolding and capacity building [22, 25, 49, 50, 78]. More importantly, we note that individual judgment about AI outputs is fragile, especially as systems evolve rapidly. But when communities develop shared heuristics for engaging with the system or evaluating its outputs,

they build collective adaptive capacity that persists beyond any single tool. AI literacy becomes a community resource, transforming individual vulnerability into collective resilience [25]. This aligns with prior HCI research. For instance, Kotturi et al. demonstrate that introducing AI tools in community contexts can reduce fear of being left behind and create opportunities for collective learning [48], while other work shows entrepreneurs socially negotiate norms around AI adoption, underscoring that decisions about what counts as appropriate use are rarely made in isolation [78]. We extend this by arguing that systems should actively support these communal processes. In the context of AI-supported business planning, this means treating AI literacy as socially constructed and context-specific, rather than as a fixed curriculum to be delivered to individuals.

Treating AI literacy as socially constructed also means recognizing that communities must be able to collectively choose not to adopt, that is, to refuse. Feminist scholarship on technology refusal, such as Manifest-No (a collective statement against extractive AI systems [59]) and SUPERRR Lab's statement on AI (articulating concerns from intersectional feminist and ecological justice perspectives [86]), positions collective refusal as an act of agency and political resistance to harmful technologies, extending HCI's view of refusal as an informed stance when technologies misalign with one's goals, values, or constraints [11, 27, 80, 81]. In our study, refusal operated at two levels: formal refusal to participate in data collection (only 30 of 48 users opted in) and practical refusal of AI use for specific tasks. P9 exemplified the latter, describing herself as "not a fan of AI" for "political, environmental, [and] creative" reasons and having "never even opened ChatGPT." However, our methodological choice to de-center AI in workshop framing may have inadvertently made informed refusal harder—by not explicitly naming AI's role upfront, we risked obscuring what participants were choosing to engage with or refuse. This tension highlights that designing for resilience includes designing for informed refusal by making the stakes of AI adoption visible enough for communities to critically evaluate when and how these tools fit their needs [48].

BizCHAT operationalized these principles by positioning business plans as drafts, encouraging participants to share, revise, and reflect with others—turning AI use into a site of communal sense-making where adoption, adaptation, and refusal could all be negotiated collectively. Thus, we call for designing AI systems that intentionally scaffold community-based resilience practices: by positioning outputs as negotiable, making collaboration and feedback around AI-generated content easy, and creating affordances for groups (not just individuals) to build shared capacity and literacy together.

However, enabling these communal scaffolds is not enough on its own. They must also be supported by systems that communities can appropriate and use on their own terms. Without control over how the system fits into their practices, communities may struggle to sustain engagement or build confidence in the tool. Thus, we turn next to the challenge of designing for co-optability, exploring how systems can support forms of ownership beyond technical maintenance and enable communities to take ownership of the processes that shape their use.

6.1.3 Co-optability. In our deployment, P4 and P7 expressed desire to bring BizCHAT into their own communities; P7 to his library and entrepreneurial hub, P4 to fellow entrepreneurs she mentors. They valued the system. However, neither wanted to simply reuse the tool. They wanted control over how it was introduced, who was in the room, and how it was framed. At the same time, they were aware of the gap in their technical skills and infrastructural resources to run and maintain it. This reflects a form of ownership over the process, one that goes beyond direct platform stewardship. We ask, how can we support the community to have such ownership?

We argue for emphasizing *ownership of processes over ownership of the product*. Communities should not be expected to host servers or manage GitHub issues if these are not relevant goals for them, but they can exercise ownership in how the tools are introduced, adapted, and integrated locally. P7's comment that "*if I did it, they'd believe it*" illustrates one way process ownership can be tied to messenger credibility and trust; who delivers the technology matters. He likened introducing new tools to "*going to get the bread*," which, in his framing, positioned knowledge-sharing as carrying risk and responsibility, and suggested that a trusted messenger may be key to ensuring the message is received. If systems cannot be reframed and delivered by these local messengers, they may risk being perceived as "*from across the tracks*" and, as a result, struggle to take root. This matters for resilience because tools introduced by outsiders risk being abandoned when external support ends. When communities own the process of introduction and adaptation, they build capacity to sustain and evolve their use over time, even as specific tools change [49].

Designing for co-optability, then, means creating systems that make it easy for communities to take and make the tool their own. In practice for BizCHAT and other generative AI tools, this could include lightweight administrator roles that allow local facilitators to configure onboarding questions, adjust framing to match local language, and provide just-in-time support for "*laundry list*" operational skills (e.g., password management, account setup) that often stand as hidden barriers to adoption. Beyond technical access, systems can offer ways for facilitators to submit feature requests, bug reports, and workshop feedback that directly integrate into maintainers' workflows, enabling a model where the community helps shape how the system is maintained and how maintenance is prioritized. Future work could explore interfaces that allow community facilitators to configure model selection and communicate tradeoffs, such as privacy, cost, and accuracy, to participants.

6.2 Methodological Implications: Deployment Details and Reflections

We chose business planning as a context for resilience building because it is a practice that facilitates reflection, can yield tangible results (e.g., access to capital), and where entrepreneurs can have expertise over the outputs (see Section 3.1). We explicitly framed BizCHAT's outputs as drafts rather than finished products to encourage help-seeking—in service of our broader framing of supporting resilience. Beyond the tool's design, our deployment methodology was also aligned with a resilience orientation. Given that accountability is an ongoing process, and following critical work in HCI which emphasizes participant-researcher reciprocity, we explicitly

positioned follow-up interviews as accountability check-ins for entrepreneurs, in service of progress on their business plan.

Another way we foregrounded resilience was by explicitly advertising and framing the workshops as “Resilience and Uncertainty Planning Workshops,” making participants’ business goals the focal point of the sessions, rather than the presence of AI. We adopted this framing as a way to create an accessible entry point into the workshops—participants were not expected to already be comfortable with AI—and to avoid attracting those only motivated by the novelty of generative AI. This decision succeeded in reaching entrepreneurs focused on business planning.

However, it also introduced tensions around transparency. Because we de-emphasized AI in our outreach materials, several participants shared that it was not immediately clear that AI would be part of the workshop. One participant (P9) explicitly raised concerns about this lack of transparency, highlighting a dangerous tension: our framing, while inviting broader participation, risked eroding trust with community collaborators. Future deployments should make AI’s role explicit upfront, even when de-centering it pedagogically.

Given that many commercial AI platforms retain user data to train models, our opaqueness around the use of AI also conflicted with IP concerns among entrepreneurs. These concerns created an important methodological and operational constraint for our deployment, motivating us to reflect carefully on which models we used and why. Throughout the two-year development of BizCHAT, there were no standardized APIs for LLMs. While we initially implemented on top of open-source models and a private server, their rapid improvement cycles made it difficult for our research team to keep pace, and hosting them at comparable latency to commercial APIs was costly. Given the small scale and local nature of our deployment, it was not operationally feasible to self-host models. Today, with the emergence of standardized API specifications (e.g., the OpenAI SDK spec [68]) and routing tools like OpenRouter [3] that allow users to select models through a common interface, there is a more feasible path to balancing this tension: supporting rapid iteration on a small team while reducing dependence on state-of-the-art commercial providers. However, each platform still maintains its own data retention and model training policies, creating uncertainty for participants and also for us. For small research teams, even hosting their own infrastructure for open-source models can be prohibitively expensive—both in compute costs and in the expertise required to maintain production-level reliability. We argue that this highlights a broader infrastructural gap in HCI research: there are few shared, privacy-preserving resources that enable researchers to deploy cutting-edge models without defaulting to commercial APIs. Creating such a resource pool would lower the barrier for small teams, making it possible to both serve state-of-the-art models and protect participants’ data. As more researchers investigate small-scale deployments, where the cost and infrastructural trade-offs of serving open-source models are prohibitive, addressing this gap could make experimentation more equitable, reproducible, and accessible. In practice, this could be implemented on top of existing research infrastructures like NSF CloudBank [2]. Building from *productive friction* and *communal scaffolds for AI literacy*, in future deployments of BizCHAT, as we introduce more open source

models, productive friction around model selection could be another point where communities negotiate tradeoffs and support AI literacy.

6.3 Limitations

There were several limitations to this study. First, we did not conduct formal evaluations of business plan quality, partly due to leadership turnover among Prototype staff. While we mitigated this through ongoing accountability check-ins, this limited our ability to assess long-term impact and formally evaluate the generated business plans. Future work can develop evaluation approaches that assess whether AI systems support entrepreneurial agency rather than merely producing professionally formatted documents. For instance, one promising direction is to evaluate whether key business details that entrepreneurs articulate during onboarding chat conversations are preserved rather than summarized out—a frequent occurrence [33]—in the final generated plan. Such an approach would assess not just output quality, but whether the system successfully supports entrepreneurs in maintaining agency and specificity over their business narrative. Ultimately, as business plans serve as both a reflective scaffold and document to access capital, longer-term evaluations could measure how well business plans supported these two ends as they are relevant to a particular BizChat user’s goals.

Second, as with many community-centered deployments, our findings are shaped by the specific context and relationships in which BizCHAT was embedded. The socio-economic dynamics, including community-level trust and support structures for entrepreneurs, can differ in other contexts. Additionally, while participants such as P4 and P7 expressed intentions to act as trusted messengers introducing BizCHAT within their communities, this study did not observe whether such introductions led to subsequent community adoption or sustained AI use. Future work should empirically examine how messenger credibility, local trust, and community identity shape whether AI tools are taken up, adapted, or refused over time.

Third, our interview sample size was small (10 entrepreneurs and support personnel), and participation was opt-in, introducing self-selection bias that likely reflects who felt empowered or resourced to engage with an AI tool. Finally, survey completion rates were low, limiting our ability to draw quantitative conclusions about changes in entrepreneurial confidence or AI comfort—though qualitative interviews provided richer insights into these dynamics. In future work, we will investigate balancing the need for survey completion with reducing extractive data collection practices, such as by communicating more clearly to users the purpose of surveys for improvement of the system and support of entrepreneurs more broadly. Thus, the findings reported here may not be generalizable; however, the lessons learned from the design and deployment can be transferable to other contexts and inform researchers interested in working with entrepreneurs or in fostering resilience amid rapid technological change.

7 Conclusion

In this paper, we introduced BizCHAT to a feminist makerspace and entrepreneurial hub in Pittsburgh, Pennsylvania. In doing so,

we moved beyond asking whether generative AI belongs in entrepreneurial practice to exploring what happens when it is woven into existing business and communal workflows. We showed how a community-centered design and deployment approach not only lowered technical barriers but also strengthened the social scaffolds that enable entrepreneurs to build resilience. Our findings revealed a central tension: while BizCHAT reduced barriers to accessing capital by translating ideas into "business language," this ease raised questions about whether instant AI outputs undermine the reflective sensemaking essential to planning. Entrepreneurs navigated this tension collectively—transforming uncertainty about AI into shared strategies for evaluating, contextualizing, and when appropriate, refusing outputs. AI is not the first technology to exacerbate inequalities in adoption, and it certainly will not be the last. In designing for resilience, we argue that the real challenge is designing through the hype toward practical, situated use. Amid the hype of each new innovation, we are tempted to ask how our practices around design should change. But as our findings show, it is human resilience and community infrastructure that persist, and our task as designers is to respect and reinforce those systems.

Acknowledgments

First and foremost, we thank our participants and the Prototype PGH community. This work was funded, in part, by UMBC's Alex. Brown Center for Entrepreneurship. Generative AI tools assisted in the preparation of this work: ChatGPT (GPT 5.1) and Claude (Sonnet 4.5) were used for summarization and refinement of text; Gemini's Nano Banana was used to generate assets for Figure 1.

A Survey Questions and Results

Participants responded to a pre-survey ($n=10$) and post-survey ($n=4$) measuring clarity, uncertainty, and AI comfort. Average Likert-scale scores (1–5) are reported below with change (Δ) computed as Post – Pre.

- **Clarity on Steps:** How clear are you on what steps your business needs to take to move forward? *Pre:* 3.32, *Post:* 3.25 ($\Delta = -0.07$)
- **Uncertainty (Strategic):** I am uncertain about the *strategic direction* of my business. *Pre:* 3.34, *Post:* 3.00 ($\Delta = -0.34$)
- **Uncertainty (Technical):** I am uncertain about the *technical aspects* of my business. *Pre:* 3.18, *Post:* 3.00 ($\Delta = -0.18$)
- **Uncertainty (Marketing):** I am uncertain about the *marketing aspects* of my business. *Pre:* 2.87, *Post:* 3.00 ($\Delta = +0.13$)
- **Uncertainty (Operations):** I am uncertain about managing the *operations* of my business. *Pre:* 3.96, *Post:* 3.00 ($\Delta = -0.96$)
- **Uncertainty (Support):** I am uncertain where to seek trustworthy *social support and guidance*. *Pre:* 3.81, *Post:* 3.00 ($\Delta = -0.81$)
- **AI Comfort:** I feel comfortable using AI tools in my business planning. *Pre:* 3.54, *Post:* 3.00 ($\Delta = -0.54$)

References

- [1] Upmetrics 2025. *AI-Powered Upmetrics Assistant*. Upmetrics. <https://upmetrics.co/features/ai-assistant>
- [2] 2025. CloudBank: Cloud Computing for Research & Education. <https://www.cloudbank.org/>. Accessed: 2025-09-11.
- [3] 2025. OpenRouter: The Unified Interface for LLMs. <https://openrouter.ai>. Accessed: 2025-09-11.
- [4] SCORE Association 2025. *SCORE: For the Life of Your Business*. SCORE Association. <https://www.score.org/>
- [5] IdeaBuddy 2025. *Smart business plan builder, powered by AI*. IdeaBuddy. <https://ideabuddy.com/features/business-plan/>
- [6] Jacob Abbott, Haley MacLeod, Novia Nurain, Gustave Ekobe, and Sameer Patil. 2019. Local standards for anonymization practices in health, wellness, accessibility, and aging research at CHI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [7] U.S. Small Business Administration. 2025. Write Your Business Plan: Sample Business Plans & Business Plan Template. <https://www.sba.gov/business-guide/plan-your-business/write-your-business-plan>. Accessed: 2025-09-11.
- [8] Sherry R Arnstein. 1969. A ladder of citizen participation. *Journal of the American Institute of planners* 35, 4 (1969), 216–224.
- [9] Seyram Avle, Julie Hui, Silvia Lindtner, and Tawanna Dillahunt. 2019. Additional labors of the entrepreneurial self. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [10] François Bar, Francis Pisani, and Matthew Weber. 2007. Mobile technology appropriation in a distant mirror: Baroque infiltration, creolization and cannibalism. *Seminario sobre Desarrollo Económico, Desarrollo Social y Comunicaciones Móviles en América Latina* (2007), 20–21.
- [11] Eric PS Baumer, Morgan G Ames, Jed R Brubaker, Jenna Burrell, and Paul Dourish. 2014. Refusing, limiting, departing: why we should study technology non-use. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. 65–68.
- [12] Ruha Benjamin. 2016. Informed refusal: Toward a justice-based bioethics. *Science, Technology, & Human Values* 41, 6 (2016), 967–990.
- [13] Rahul Bhargava, Erica Deahl, Emmanuel Letouzé, Amanda Noonan, David Sangokoya, and Natalie Shoup. 2015. Beyond data literacy: Reinventing community engagement and empowerment in the age of data. (2015).
- [14] Abeba Birhane, William Isaac, Vinodkumar Prabhakaran, Mark Diaz, Madeleine Clare Elish, Jason Gabriel, and Shakir Mohamed. 2022. Power to the people? Opportunities and challenges for participatory AI. In *Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–8.
- [15] George A. Bonanno. 2004. Loss, Trauma, and Human Resilience: Have We Underestimated the Human Capacity to Thrive After Extremely Aversive Events? *American Psychologist* 59, 1 (2004), 20–28. <https://doi.org/10.1037/0003-066X.59.1.20>
- [16] Callum Borchers. [n. d.]. Americans Are Side-Hustling Like We're in a Recession. *The Wall Street Journal* [n. d.]. Online; accessed 2025-09-07.
- [17] K.J. Bucknell, M. Kangas, and M.F. Crane. 2022. Adaptive self-reflection and resilience: The moderating effects of rumination on insight as a mediator. *Personality and Individual Differences* 185 (2022), 111234. <https://doi.org/10.1016/j.paid.2021.111234>
- [18] Pew Research Center. 2023. *What the Data Says About Americans' Views of Artificial Intelligence*. <https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/> Accessed: 2025-01-07.
- [19] Shannon Cook. 2021. The Rise Of TikTok Entrepreneurs: How Useful Is Their Business Advice? *BusinessBecause* (27 May 2021). <https://www.businessbecause.com/news/insights/7658/tiktok-entrepreneurs-business-advice> Accessed: 2025-09-11.
- [20] Ned Cooper, Tiffanie Horne, Gillian R Hayes, Courtney Heldreth, Michal Lahav, Jess Holbrook, and Lauren Wilcox. 2022. A systematic review and thematic analysis of community-collaborative approaches to computing research. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–18.
- [21] Fernando Delgado, Stephen Yang, Michael Madaio, and Qian Yang. 2023. The participatory turn in ai design: Theoretical foundations and the current state of practice. In *Proceedings of the 3rd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–23.
- [22] Tawanna R Dillahunt. 2014. Fostering social capital in economically distressed communities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 531–540.
- [23] Tawanna R Dillahunt, Vaishnav Kameswaran, Desiree McLain, Minnie Lester, Delores Orr, and Kentaro Toyama. 2018. Entrepreneurship and the socio-technical chasm in a lean economy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [24] Tawanna R. Dillahunt, Alex Jiahong Lu, Aarti Israni, Ruchita Lodha, Savana Brewer, Tiera S. Robinson, Angela Brown Wilson, and Earnest Wheeler. 2022. The Village: Infrastructuring Community-based Mentoring to Support Adults Experiencing Poverty. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New Orleans, LA, USA, 1–17. <https://doi.org/10.1145/3491102.3501949>
- [25] Tawanna R Dillahunt, Kerby Shedden, Mila Ekaterina Filipof, Soyoung Lee, Mustafa Naseem, Kentaro Toyama, and Julie Hui. 2025. The Development of a New Measure of Collective Digital Literacy. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–33.

- [26] John C. Flanagan. 1954. The Critical Incident Technique. *Psychological Bulletin* 51, 4 (1954), 327–358. <https://doi.org/10.1037/h0061470>
- [27] Aakash Gautam. 2024. Reconfiguring participatory design to resist ai realism. In *Proceedings of the Participatory Design Conference 2024: Exploratory Papers and Workshops-Volume 2*. 31–36.
- [28] Aakash Gautam and Deborah Tatar. 2020. p for political: Participation without agency is not enough. In *Proceedings of the 16th Participatory Design Conference 2020-Participation (s) Otherwise-Volume 2*. 45–49.
- [29] Aakash Gautam and Deborah Tatar. 2022. Empowering Participation Within Structures of Dependency. In *Proceedings of the Participatory Design Conference 2022-Volume 1*. 75–86.
- [30] Aakash Gautam, Deborah Tatar, and Steve Harrison. 2020. Crafting, communality, and computing: Building on existing strengths to support a vulnerable population. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [31] Elena L. Glassman, Ziwei Gu, and Jonathan K. Kummerfeld. 2024. Ai-resilient interfaces. *arXiv preprint arXiv:2405.08447* (2024).
- [32] Peter M. Gollwitzer. 1999. Implementation intentions: strong effects of simple plans. *American psychologist* 54, 7 (1999), 493.
- [33] Tanya Goyal, Junyi Jessy Li, and Greg Durrett. 2022. News summarization and evaluation in the era of gpt-3. *arXiv preprint arXiv:2209.12356* (2022).
- [34] Gusto. 2024. *How SMBs Are Using AI in 2024*. <https://gusto.com/company-news/smb-using-ai-2024/> Accessed: 2025-01-07.
- [35] Luke Haliburton, David Joachim Grüning, Frederik Riedel, Albrecht Schmidt, and Nada Terzimehić. 2024. A longitudinal in-the-wild investigation of design frictions to prevent smartphone overuse. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [36] Christina Harrington, Sheena Erete, and Anne Marie Piper. 2019. Deconstructing community-based collaborative design: Towards more equitable participatory design engagements. *Proceedings of the ACM on human-computer interaction* 3, CSCW (2019), 1–25.
- [37] Suzanne Hidi and K. Ann Renninger. 2006. The four-phase model of interest development. *Educational psychologist* 41, 2 (2006), 111–127.
- [38] Theo Hug. 2005. Micro Learning and Narration. Exploring possibilities of utilization of narrations and storytelling for the designing of "micro units" and didactical micro-learning arrangements. In *fourth Media in Transition conference*, Vol. 6. Mit4.
- [39] Julie Hui, Nefer Ra Barber, Wendy Casey, Suzanne Cleage, Danny C. Dolley, Frances Worthy, Kentaro Toyama, and Tawanna R. Dillahunt. 2020. Community collectives: Low-tech social support for digitally-engaged entrepreneurship. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–15.
- [40] Julie Hui, Kristin Seefeldt, Christie Baer, Lutalo Sanifu, Aaron Jackson, and Tawanna R. Dillahunt. 2023. Community Tech Workers: Scaffolding Digital Engagement Among Underserved Minority Businesses. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–25.
- [41] Julie Hui, Kentaro Toyama, Joyojeet Pal, and Tawanna Dillahunt. 2018. Making a living my way: Necessity-driven entrepreneurship in resource-constrained communities. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–24.
- [42] Chris Iacofano. 2025. *LivePlan Assistant*. Palo Alto Software. <https://help.liveplan.com/hc/en-us/articles/1492589421671-LivePlan-Assistant>
- [43] Sarah Inman and David Ribes. 2019. "Beautiful Seams" Strategic Revelations and Concealments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [44] Barbara A. Israel, Amy J. Schulz, Edith A. Parker, and Adam B. Becker. 1998. Review of Community-Based Research: Assessing Partnership Approaches to Improve Public Health. *Annual Review of Public Health* 19 (1998), 173–202. <https://doi.org/10.1146/annurev.publhealth.19.1.173>
- [45] Philip Jackson. 2002. Life in classrooms. In *Teaching and Learning in the Primary School*. Routledge, 123–128.
- [46] Manu Kapur. 2008. Productive failure. *Cognition and instruction* 26, 3 (2008), 379–424.
- [47] Louise Kay and Amy Buxton. 2023. Makerspaces and the Characteristics of Effective Learning in the Early Years. *Journal of Early Childhood Research* 22, 3 (2023), 343–358. <https://doi.org/10.1177/1476718X231210633> Original work published 2024.
- [48] Yasmine Kotturi, Angel Anderson, Glenn Ford, Michael Skirpan, and Jeffrey P. Bigham. 2024. Deconstructing the Veneer of Simplicity: Co-Designing Introductory Generative AI Workshops with Local Entrepreneurs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–16.
- [49] Yasmine Kotturi, Julie Hui, TJ Johnson, Lutalo Sanifu, and Tawanna R. Dillahunt. 2024. Sustaining Community-Based Research in Computing: Lessons from Two Tech Capacity Building Initiatives for Local Businesses. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–31.
- [50] Yasmine Kotturi, Herman T. Johnson, Michael Skirpan, Sarah E. Fox, Jeffrey P. Bigham, and Amy Pavel. 2022. Tech help desk: Support for local entrepreneurs addressing the Long Tail of computing challenges. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [51] Yasmine Kotturi, Jenny Yu, Pranav Khadpe, Erin Gatz, Harvey Zheng, Sarah E. Fox, and Chinmay Kulkarni. 2024. Peerdeas: Co-Designing a Peer Support Platform with Creative Entrepreneurs. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–24.
- [52] Kristine M. Kuhn and Tera L. Galloway. 2015. With a Little Help from My Competitors: Peer Networking among Artisan Entrepreneurs. *Entrepreneurship Theory and Practice* 39, 3 (2015), 571–600. <https://doi.org/10.1111/etap.12053>
- [53] Johan P. Larsson and Per Thulin. 2019. Independent by necessity? The life satisfaction of necessity and opportunity entrepreneurs in 70 countries. *Small Business Economics* 53, 4 (2019), 921–934.
- [54] Rebecca Leppert. 2024. *A look at small businesses in the U.S.* Accessed 2025-09-07.
- [55] Keyao Li, Mark A. Griffin, and Mengting (Rachel) Xia. 2025. How do work-force adaptability and reskilling initiatives drive innovations: the case of Western Australian construction industry. *Construction Management and Economics* 43, 9 (2025), 746–763. <https://doi.org/10.1080/01446193.2025.2511831>
- [56] Calvin A. Liang, Sean A. Munson, and Julie A. Kientz. 2021. Embracing four tensions in human-computer interaction research with marginalized people. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 2 (2021), 1–47.
- [57] Duri Long and Brian Magerko. 2020. What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–16.
- [58] Suniya S. Luthar. 1991. Vulnerability and Resilience: A Study of High-Risk Adolescents. *Child Development* 62, 3 (1991), 600–616. <https://doi.org/10.1111/j.1467-8624.1991.tb01555.x>
- [59] Manifest-No Collective. 2024. Manifest-No. <https://www.manifestno.com/>. Accessed: 2025-12-01.
- [60] Ann S. Masten. 2001. Ordinary magic: Resilience processes in development. *American psychologist* 56, 3 (2001), 227.
- [61] Microsoft. 2025. Prompts to try. <https://support.microsoft.com/en-us/copilot-microsoft365-chat> Accessed: 2025-04-27.
- [62] Meredith Ringel Morris, Carrie J. Cai, Jess Holbrook, Chinmay Kulkarni, and Michael Terry. 2023. The Design Space of Generative Models. *arXiv preprint arXiv:2304.10547* (2023). <https://arxiv.org/abs/2304.10547> Accessed: 2025-09-11.
- [63] Helene Mueller and Martina Pieperhoff. 2023. Necessity entrepreneurship: An integrative review and research agenda. *Entrepreneurship & Regional Development* 35, 9–10 (2023), 762–787.
- [64] Michael Muller and Allison Druin. 2002. Participatory Design: The Third Space in HCI. *Handbook of HCI* (01 2002).
- [65] Diana Murdoch, Andrea R. English, Allison Hintz, and Kersti Tyson. 2020. Feeling heard: Inclusive education, transformative learning, and productive struggle. *Educational theory* 70, 5 (2020), 653–679.
- [66] Fran H. Norris, Susan P. Stevens, Betty Pfefferbaum, Karen F. Wyche, and Rose L. Pfefferbaum. 2007. Community Resilience as a Metaphor, Theory, Set of Capacities, and Strategy for Disaster Readiness. *American Journal of Community Psychology* 41 (2007), 127–150. <https://doi.org/10.1007/s10464-007-9156-6>
- [67] OpenAI. 2025. ChatGPT Overview. <https://openai.com/chatgpt/overview/> Accessed: 2025-01-03.
- [68] OpenAI. 2025. openai/openai-openapi: OpenAPI specification for the OpenAI API. <https://github.com/openai/openai-openapi>. Accessed: 2025-09-11.
- [69] Nicholas Otis, Rowan P. Clarke, Solene Delecourt, David Holtz, and Rembrand Koning. 2023. The uneven impact of generative AI on entrepreneurial performance. *Available at SSRN* 4671369 (2023).
- [70] Seymour A. Papert. 2020. *Mindstorms: Children, computers, and powerful ideas*. Basic books.
- [71] Andrea G. Parker, Laura M. Vardoulakis, Jatin Alla, and Christina N. Harrington. 2025. Participatory AI Considerations for Advancing Racial Health Equity. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–24.
- [72] Shantanu S. Patel, Monica S. Rogers, Jason Amlôt, and G. James Rubin. 2017. What Do We Mean by "Community Resilience"? A Systematic Literature Review of How It Is Defined in the Literature. *PLoS Currents* 9 (2017). <https://doi.org/10.1371/currents.dis.db775aff25efc5ac4f0660ad9c9f7db2>
- [73] Lucy Pei and Bonnie Nardi. 2019. We did it right, but it was still wrong: Toward assets-based design. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [74] Jennifer Pierre, Roderic Crooks, Morgan Currie, Britt Paris, and Irene Pasquetto. 2021. Getting Ourselves Together: Data-centered participatory design research & epistemic burden. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [75] Peter Pirolli. 2007. *Information Foraging Theory: Adaptive Interaction with Information*. Oxford University Press, Oxford, UK.
- [76] Mitchel Resnick. 2008. Falling in love with Seymour's ideas. In *American Educational Research Association (AERA) annual conference*.
- [77] Jennifer R. Riley and Ann S. Masten. 2005. Resilience in context. In *Resilience in children, families, and communities: Linking context to practice and policy*. Springer, 13–25.

- [78] Quentin Romero Lauro, Jeffrey P Bigham, and Yasmine Kotturi. 2024. Exploring the Role of Social Support When Integrating Generative AI in Small Business Workflows. In *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing*. 485–492.
- [79] Michael Rutter. 2012. Resilience as a Dynamic Concept. *Development and Psychopathology* 24, 2 (2012), 335–344. <https://doi.org/10.1017/S0954579412000028>
- [80] Christine Satchell and Paul Dourish. 2009. Beyond the user: use and non-use in HCI. In *Proceedings of the 21st annual conference of the Australian computer-human interaction special interest group: Design: Open 24/7*. 9–16.
- [81] Neil Selwyn. 2006. Digital division or digital decision? A study of non-users and low-users of computers. *Poetics* 34, 4-5 (2006), 273–292.
- [82] Mrigank S Shail. 2019. Using micro-learning on mobile applications to increase knowledge retention and work performance: a review of literature. *Cureus* 11, 8 (2019).
- [83] Steven M. Southwick, George A. Bonanno, Ann S. Masten, Catherine Panter-Brick, and Rachel Yehuda. 2014. Resilience Definitions, Theory, and Challenges: Interdisciplinary Perspectives. *European Journal of Psychotraumatology* 5 (2014), 25338. <https://doi.org/10.3402/ejpt.v5.25338>
- [84] Stripe. [n. d.]. Stripe Atlas: Turn your idea into a startup. <https://stripe.com/atlas> Library Catalog: stripe.com.
- [85] Cella M Sum, Jiayin Zhi, Amil NT Cook, Patrick James Cooper, Arturo Lozano, TJ Johnson, Jason Perez, Rayid Ghani, Michael Skirpan, Motahhare Eslami, et al. 2025. "You're in a Ferrari. I'm Waiting for the Bus": Confronting Tensions in Community-University Partnerships. *Proceedings of the ACM on Human-Computer Interaction* 9, 2 (2025), 1–28.
- [86] SUPERRR Lab. 2024. About AI and Unlikelihood. <https://superrr.net/en/blog/about-ai-and-unlikelihood>. Accessed: 2025-12-01.
- [87] Catrien JAM Termeer and Margo A van den Brink. 2013. Organizational conditions for dealing with the unknown unknown: Illustrated by how a Dutch water management authority is preparing for climate change. *Public Management Review* 15, 1 (2013), 43–62.
- [88] Michael Ungar. 2013. Resilience, Trauma, Context, and Culture. *Trauma, Violence, & Abuse* 14, 3 (2013), 255–266. <https://doi.org/10.1177/1524838013487805>
- [89] U.S. Small Business Administration. n.d. Write Your Business Plan. <https://www.sba.gov/business-guide/plan-your-business/write-your-business-plan> Accessed: 2025-01-07.
- [90] John Vines, Rachel Clarke, Peter Wright, John McCarthy, and Patrick Olivier. 2013. Configuring participation: on how we involve people in design. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 429–438.
- [91] Eric von Hippel. 2006. *Democratizing Innovation*. The MIT Press, Cambridge, MA.
- [92] Dhaval Vyas and Tawanna Dillahunt. 2017. Everyday resilience: Supporting resilient strategies among low socioeconomic status communities. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 1–21.
- [93] Hiroko Kawaguchi Warshauer. 2015. Productive struggle in middle school mathematics classrooms. *Journal of Mathematics Teacher Education* 18, 4 (2015), 375–400.
- [94] Karl E Weick and Karl E Weick. 1995. *Sensemaking in organizations*. Vol. 3. Sage publications Thousand Oaks, CA.
- [95] Marisol Wong-Villares, Aakash Gautam, Deborah Tatar, and Betsy DiSalvo. 2021. Reflections on assets-based design: A journey towards a collective of assets-based thinkers. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–32.
- [96] JD Zamfirescu-Pereira, Richmond Y Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–21.