INTERNATIONAL JOURNAL OF ACADEMIC RESEARCH IN BUSINESS & SOCIAL SCIENCES

Vol 16, Issue 11, (2026) E-ISSN: 2222-6990

AI-Mediated Care-Washing: Discourse, Cognition, and Dual-Layer Decoupling in China's Elderly Companion Product Sector

Xu Jia, Noor Fareen Abdul Rahim*

Graduate School of Business, Universiti Sains Malaysia, Penang, Malaysia

Email: xujia1999@student.usm.my

*Corresponding Author Email: noorfareen@usm.my

DOI Link: http://dx.doi.org/10.6007/IJARBSS/v16-i11/29421

Published Date: 25 November 2026

Abstract

This study examines how enterprises in China's AI companion product sector publicly construct their products as legitimate care solutions for elderly users, and whether practitioners privately understand that discourse's relationship to commercial and technical realities. A two-phase sequential qualitative design was employed. Study 1 conducted qualitative content analysis of twelve enterprise-affiliated texts through hybrid deductive-inductive coding; Study 2 conducted semi-structured interviews with eight practitioners in mainland China and Hong Kong, analyzed through thematic analysis. Study 1 identified two discourse patterns: emotional equivalence claims constructing AI interaction as equivalent to human companionship, and vulnerability erasure replacing elderly vulnerability with technological empowerment. Study 2 found that practitioners acknowledged a gap between promotional claims and technical capability, maintained different communicative registers for internal and external audiences, and were privately aware of the vulnerability enterprise discourse suppresses. This study proposes AI-mediated care-washing as a dual-layer decoupling construct extending care-washing theory, contributes to institutional decoupling theory by documenting communicative register bifurcation, and advances vulnerable consumer research by identifying filial obligation as the mechanism converting elderly vulnerability into purchasing behaviour.

Keywords: Care-Washing, Ai Companion Products, Silver Economy, Institutional Decoupling, Vulnerable Consumers, Legitimacy Theory

Introduction

China now has over 300 million citizens aged 60 and above, and its eldercare workforce cannot meet projected demand. In this context, AI-powered companion products have emerged as a commercially attractive and policy-endorsed response (National Health Commission of China, 2022). Enterprises have positioned these products as solutions to elderly loneliness and absent family presence, claiming capacities for emotional support and

social substitution that current technology does not demonstrably possess (Pu et al., 2025). These claims are striking not because they are obviously false, but because they are largely unverifiable: no agreed standard exists for what it means for a machine to provide care, and no regulatory mechanism currently requires enterprises to substantiate the emotional outcomes they promise. The result is a promotional environment in which care language proliferates without accountability.

Existing research has examined AI companion technology from the perspectives of user acceptance (Cesta et al., 2016), psychological outcomes (Abdi et al., 2018), and care ethics (Sharkey and Sharkey, 2012; Vallor, 2011), but has not analysed the promotional discourse through which enterprises construct these products as legitimate care solutions. Chatzidakis and Littler (2022) introduced care-washing to describe the strategic commodification of care through corporate branding, but their framework addressed conventional consumer goods and did not account for the additional misrepresentation involved when a product claims emotional outcomes that the underlying technology cannot demonstrably deliver. To analyse this phenomenon, the study draws on legitimacy theory (Suchman, 1995), which explains why enterprises strategically construct care-oriented discourse to secure moral approval in socially sensitive markets, and institutional decoupling theory (Meyer and Rowan, 1977), which accounts for the structural separation between formal promotional discourse and actual operational logic. This gap motivates the present study.

This study employs a two-phase qualitative design. Study 1 conducted content analysis of twelve enterprise-affiliated texts from China's AI companion and smart ageing product sector, identifying two primary discourse patterns: emotional equivalence claims, in which AI interaction is constructed as equivalent to human companionship, and vulnerability erasure, in which elderly cognitive and social vulnerability is replaced by a discourse of technological empowerment. Study 2 conducted semi-structured interviews with eight industry practitioners in mainland China and Hong Kong, analysed through thematic analysis following Braun and Clarke (2006). Practitioners universally acknowledged a gap between promotional claims and technical capability, maintained different communicative registers for internal and external audiences, and were privately aware of the elderly vulnerability that enterprise discourse suppresses.

The study makes three contributions. It extends care-washing theory by identifying a dual-layer decoupling structure specific to AI contexts, in which commercial logic is concealed behind care language and technical uncertainty is concealed behind capability claims. This construct is designated AI-mediated care-washing. It contributes to institutional decoupling theory (Meyer and Rowan, 1977) by documenting communicative register bifurcation as the practitioner-level mechanism sustaining that structure. And it contributes to vulnerable consumer research (Baker et al., 2005) by demonstrating that the commercial exploitation of elderly vulnerability operates through its discursive suppression, with filial obligation serving as the mechanism through which care language converts family emotional motivation into purchasing behaviour.

Literature Review

Legitimacy Theory and the Strategic Construction of Care Discourse

The analytical framework required to examine enterprise promotional discourse in a care-sensitive market begins with legitimacy theory (Suchman, 1995). Suchman (1995) distinguishes three forms of organisational legitimacy: pragmatic legitimacy, grounded in the immediate self-interest of audiences; moral legitimacy, grounded in normative evaluations of whether the organisation is doing the right thing; and cognitive legitimacy, grounded in the taken-for-granted assumption that the organisation's activities are natural and necessary. Enterprises operating in markets associated with vulnerable populations are subject to heightened moral legitimacy demands: their commercial activity must be publicly justified not only as effective but as ethically appropriate. Promotional discourse is a primary site through which this justification is produced and sustained (Palazzo and Scherer, 2006; Ashforth and Gibbs, 1990).

The concept of greenwashing, introduced to describe the strategic use of environmental claims to acquire moral legitimacy without substantive environmental performance (Ramus and Montiel, 2005), has been extensively theorised as an application of legitimacy-seeking behaviour to a socially sensitive domain (Lyon and Maxwell, 2011; Delmas and Burbano, 2011; Laufer, 2003). Greenwashing research has identified a range of discursive strategies through which enterprises construct an appearance of environmental responsibility, including vague or unsubstantiated claims, selective disclosure of positive information, and the framing of commercial benefits as environmental goods (Walker and Wan, 2012; Marquis et al., 2016). The analytical logic of greenwashing has subsequently been extended to other domains of social value, including pinkwashing, in which enterprises deploy LGBTQ-inclusive imagery to acquire social legitimacy, and purpose-washing, in which social mission language is used to reframe commercial expansion as social progress (Hanleybrown et al., 2012).

Chatzidakis and Littler (2022) introduced care-washing as the application of this logic to the domain of care, defining it as the strategic commodification of care through corporate branding. Drawing on cases from the COVID-19 pandemic, they identified three operational tendencies: opportunistic branding, community resourcing, and reputational steamrolling. Their framework established care-washing as a theoretically distinct form of legitimacy-seeking behaviour in which the moral weight of care norms is appropriated to justify commercial activity. However, their analysis addressed conventional consumer goods in a general market context and did not examine AI-specific promotional dynamics. When a company sells an AI companion product while claiming it can alleviate elderly loneliness, the gap takes on an additional dimension: whether the technology can deliver the emotional outcomes being claimed remains empirically contested, and no standardised evidentiary basis currently exists against which such claims can be evaluated (Abdi et al., 2018; Sharkey and Sharkey, 2012). This technology-specific layer of misrepresentation, in which capability claims exceed demonstrable technical performance, is not addressed by the existing care-washing framework and constitutes the primary theoretical gap this study addresses.

Institutional Decoupling and the Gap Between Discourse and Practice

Legitimacy theory explains why enterprises construct care-oriented promotional discourse (Palazzo and Scherer, 2006); institutional decoupling theory explains how such discourse is structurally maintained even when those within the organisation know it does not accurately

represent their products' capabilities (Bromley and Powell, 2012). Meyer and Rowan (1977) argued that organisations operating in institutionally complex environments tend to adopt formal structures and public accounts that conform to prevailing social norms while insulating their actual technical operations from external evaluation. This decoupling of formal structure from operational activity is not a matter of individual deception but a systemic response to the contradictory demands placed on organisations by multiple institutional audiences(Bromley and Powell, 2012). DiMaggio and Powell (1983) extended this account by identifying the coercive, mimetic, and normative pressures through which organisations come to adopt similar formal structures, producing field-level homogeneity in public discourse even where operational practices vary considerably.

Applied to enterprise communication, decoupling theory predicts that organisations facing simultaneous commercial and moral legitimacy pressures will maintain systematically different communicative registers for different audiences (Cho and Patten, 2007). Research on corporate sustainability disclosure has documented this bifurcation extensively (Boiral, 2013), finding that enterprises routinely produce public communications that emphasise social and environmental commitments while internal communications prioritise financial performance, cost management, and market competitiveness (Cho et al., 2012; Boiral, 2013). In the context of AI companion products targeting elderly users, decoupling theory predicts that the care discourse documented in enterprise promotional materials will coexist with an internal operational logic oriented primarily toward cost reduction, scalability, and commercial return, and that practitioners will actively manage the boundary between these two registers depending on their intended audience.

Despite its analytical power, the application of institutional decoupling theory to enterprise communication has two limitations that the present study addresses. First, existing decoupling research has focused primarily on formal organisational documents such as sustainability reports and policy statements as the site of analysis, leaving underexamined the practitioner-level cognitive processes through which decoupling is enacted in everyday communicative practice (Boiral, 2013; Cho et al., 2012). Whether practitioners who produce decoupled discourse are aware of the gap between their public accounts and operational realities, and how they account for this gap subjectively, remains largely unaddressed. Second, decoupling research has been conducted in contexts where the performance being decoupled from is, in principle, measurable, such as carbon emissions or safety records (Marquis et al., 2016; Lyon and Maxwell, 2011). The AI companion product context introduces a different condition: the emotional outcomes being claimed are not only unverified but may be fundamentally difficult to verify with existing measurement frameworks. Existing decoupling theory does not account for this epistemic condition.

Vulnerable Consumers and the Ethical Stakes of AI Companion Discourse

The significance of care-washing in the AI companion product sector is amplified by the specific characteristics of the target population (Moschis, 2011). Baker et al. (2005) define consumer vulnerability as a state in which individuals are at heightened risk of harm from market interactions due to cognitive, physical, social, or economic conditions that limit their capacity to protect their own interests. Elderly consumers represent a population in which multiple vulnerability dimensions frequently converge: cognitive decline may impair the evaluation of complex product claims (Drolet et al., 2019); social isolation may intensify

susceptibility to products promising companionship (Alves and Wilson, 2008); limited digital literacy may prevent meaningful comparison of alternatives (Czaja et al., 2006); and fixed or declining incomes may constrain the capacity to recover from poor purchasing decisions (Moschis, 2003). These conditions do not render elderly consumers passive or incapable, but they do create conditions in which misleading promotional discourse carries greater potential for harm than in markets serving less vulnerable populations.

The ethical dimensions of deploying AI systems in care relationships for elderly users have been examined from several perspectives. Sharkey and Sharkey (2012) raised concerns about the substitution of robotic and AI interaction for human relationships in eldercare, arguing that this substitution risks diminishing the dignity and relational quality of care even when it improves efficiency. Vallor (2011) argued that care relationships depend on moral qualities, including attentiveness, empathy, and responsiveness, that AI systems cannot replicate, and that framing AI as a care provider misrepresents both the nature of care and the capabilities of the technology. Abdi et al. (2018) reviewed the evidence base for AI and robotic companion products targeting elderly users and found that while some products demonstrated benefits in specific functional domains, evidence for sustained emotional and relational outcomes remained limited and methodologically contested. These contributions establish the ethical stakes of AI companion product claims but do not analyse the promotional discourse through which those claims are constructed and sustained. The present study addresses this gap by examining how enterprises strategically deploy care language to construct AI companion products as legitimate responses to elderly loneliness, how this discourse conceals both commercial motivation and technical uncertainty, and how practitioners within the industry privately understand the relationship between the care claims they produce and the products those claims describe.

Based on the theoretical gaps and practical identified above, this study centers on the following primary research question:
RQ1: What discourse patterns emerge in the public communications of enterprises marketing AI companion products to elderly users in China?
RQ2: How do industry practitioners understand and account for the relationship between promotional discourse and the realities of AI companion products targeting elderly users?

Research Context
China's AI companion product sector constitutes an analytically distinctive research context for three reasons that collectively amplify the care-washing dynamics identified in the literature review. First, AI companion products enter a genuine structural care crisis rather than a saturated market: with a projected shortfall of approximately 10 million professional eldercare workers by 2035 and a rapidly increasing proportion of elderly people living alone or with limited family support, enterprises are positioned as apparent problem-solvers rather than market entrants, generating unusually strong conditions for moral legitimacy claims and making the boundary between authentic care provision and care-washing particularly difficult to identify and contest (Guo et al., 2019; Zhan et al., 2015). Second, this sector operates in a regulatory vacuum where no standardised definition of emotional or relational outcomes exists, no independent verification of promotional claims is required, and no enforcement mechanism exists to sanction enterprises whose products fail to deliver the companionship benefits they advertise (Roberts et al., 2021). China's regulatory frameworks for AI have

developed primarily around data security and generative AI content governance rather than the substantive verification of capability claims in consumer-facing care applications, meaning that emotional equivalence assertions can circulate without external challenge and promotional overclaiming is structurally rational from an enterprise perspective. Third, purchasing decisions for AI companion products are frequently made not by elderly users themselves but by adult children acting on Confucian norms of filial piety and guilt about physical or temporal absence from ageing parents, meaning that care-oriented promotional discourse operates on a population whose susceptibility to emotional appeals is high and whose capacity to evaluate technical claims is limited by informational asymmetry and emotional investment (Cong and Silverstein, 2008; Zhan and Montgomery, 2003). This triangulated structure is specific to the Chinese cultural context and is not captured by existing care-washing frameworks developed in Western market settings.

Overall Research Design

Guided by legitimacy theory (Suchman, 1995) and institutional decoupling theory (Meyer and Rowan, 1977), this study adopts a two-phase sequential qualitative design to theorise AI-mediated care-washing as a dual-layer discursive structure in China's AI companion product sector (Ivankova et al., 2006). Sampling across both phases is purposive and theoretically driven rather than aimed at statistical representativeness (Boiral, 2013; Braun and Clarke, 2006). Study 1 draws on a purposive corpus of twelve enterprise-affiliated texts analysed through qualitative content analysis using a hybrid deductive-inductive coding procedure, grounded in the greenwashing and organisational legitimacy literature (Fereday and Muir-Cochrane, 2006; Cho et al., 2012; Boiral, 2013). The corpus was assembled to capture the range of enterprise types active in the sector, ensuring that the identified patterns reflect a sector-level structure rather than firm-specific rhetoric. This phase produces a structured mapping of the discourse strategies through which care legitimacy is publicly constructed across the sector. Study 2 employs semi-structured interviews with eight industry practitioners in mainland China and Hong Kong, analysed through thematic analysis with an initial deductive framework derived from Study 1's discourse patterns and inductive coding conducted in parallel (Braun and Clarke, 2006). Practitioners were selected to span the functional roles most directly implicated in producing and communicating care claims. This phase reconstructs the practitioner-level cognitions and communicative register bifurcations that sustain the care-washing structure identified in Study 1, enabling the study to move from describing what enterprises say to explaining how and why they say it. Together, the two phases produce an integrated account of AI-mediated care-washing as a phenomenon that is simultaneously structural, discursive, and practitioner-enacted.

Study 1: Content Analysis of Enterprise Promotional Discourse

Corpus Construction

Study 1 examines the promotional discourse through which enterprises in China's AI companion and smart ageing product sector publicly frame their offerings. The analytical object is defined as AI-powered products whose primary marketed function is emotional companionship, social interaction, or cognitive stimulation for elderly users (Abdi et al., 2018). This definition encompasses AI digital human applications, voice-based companion assistants, and emotion-oriented smart devices, while excluding purely functional eldercare technologies such as fall-detection systems and remote health monitoring devices, which do not make companionship claims.

A purposive corpus of twelve publicly accessible enterprise-affiliated texts was constructed from Chinese-language sources published between 2023 and 2025. Four text categories were sampled to ensure coverage across the discourse ecosystem rather than reliance on any single channel: enterprise self-introductions and product promotional copy; industry research reports co-published by enterprises and research institutions; media reports containing substantial enterprise spokesperson quotations; and one critically oriented media commentary included to enable contrastive analysis. Government policy documents were retained as contextual reference only and excluded from the core coding sample. Texts were retrieved through keyword searches combining terms for silver economy, AI companionship, smart elderly care, and technology for social good. Inclusion required that texts address the companionship or emotional function of AI products for elderly users and be publicly accessible without subscription. To capture the range of enterprise types active in the sector, the sampled texts span leading platform firms, specialised AI companion startups, and traditional eldercare providers extending into AI-enabled offerings. Table 1 summarises the corpus.

Table 1
Corpus Overview (T01-T12)

Code	Source	Text Type	Year
T01	Shounanshan AI Voice Assistant	Enterprise self-introduction, industry expo;	2024
T02	Runxin Technology / Yaotong Digital	Industry media report with enterprise statement;	2025
T03	Mengyu Intelligence / Ropet	Technology media entrepreneurial report, CES product launch narrative;	2025
T04	Tencent Yuanqi Elderly Assistant	Product manager article, public interest project;	2025
T05	Tencent SSV Silver Technology Lab	Official programme launch statement;	2024
T06	Super Playmate Lulupals	Technology media promotional report;	2025
T07	Tuoshi Technology Group	Soft-news promotional article;	2023
T08	Senlikang Technology / Xiaoli companion robot	Financial/industry media feature;	2025
T09	Alibaba / Zhejiang Open University	Enterprise-academic industry research report;	2025
T10	Silver Economy Forum, South China Morning Post	Industry expert public remarks;	2024
T11	Beijing Daily	Critically oriented media commentary;	2025
T12	Xinhua News Agency	State media feature with elderly user testimony;	2025

Note. T01-T12 constitute the core coding sample.

Coding Procedure

A hybrid deductive-inductive coding approach was adopted following Fereday and Muir-Cochrane (2006). Three a priori deductive categories were derived from the greenwashing and organisational legitimacy literature prior to corpus analysis: emotional care claims asserting functional equivalence between AI interaction and human companionship; efficiency and cost logic foregrounding commercial or operational advantages as the primary value proposition; and limitation acknowledgement comprising explicit recognition that AI cannot substitute for human care (Walker and Wan, 2012; Cho and Patten, 2007). All twelve texts were read in full. Textual units were defined as individual sentences or coherent multi-sentence passages performing a single rhetorical function. Codes were applied non-exclusively. Two coders independently coded the full corpus, and coding discrepancies were resolved through discussion until consensus was reached, with the deductive framework refined iteratively where inductive patterns warranted revision. Two additional categories emerged inductively during analysis: moral alignment, in which commercial offerings are positioned within national policy discourse or public interest framing; and vulnerability erasure, in which the cognitive, social, or financial vulnerability of the elderly target population is systematically omitted from product framing. Saturation was reached after the tenth text, with no new categories emerging from T11 or T12.

Findings

Analysis of the twelve-text corpus identified two primary discourse patterns and one secondary pattern constituting the structure of AI-mediated care-washing in enterprise promotional communication. The two primary patterns, emotional equivalence claims and vulnerability erasure, are reported in detail as they directly instantiate the dual-layer decoupling structure that is the study's central theoretical contribution. The third pattern, commercial logic embedded within care framing, is reported as the structural condition that makes the first two patterns necessary. Two additional patterns, moral alignment and limitation acknowledgement, are noted briefly as contextual features.

Primary Pattern 1: Emotional Equivalence Claims

The most prevalent pattern, present in eight of twelve texts, involves the construction of AI companion products as functionally equivalent to or superior to human companionship. This pattern corresponds directly to what Suchman (1995) identifies as the pursuit of moral legitimacy: by asserting that AI interaction delivers the same relational value as human care, enterprises position themselves as ethically appropriate actors in a care-sensitive market. The Alibaba and Zhejiang Open University joint industry report (T09) provides a representative instance:

"It is always online, willing to listen, and infinitely patient. It can effectively alleviate elderly loneliness (T09)."

The attributes enumerated here, permanent availability, attentiveness, and patience, are presented as product specifications without qualification. The implicit comparison is with human caregivers, who are bounded by fatigue, absence, and competing obligations. This rhetorical move instantiates what Vallor (2011) identifies as a misrepresentation of the nature of care: the moral qualities that make care relationships meaningful, including genuine attentiveness and emotional reciprocity, are reduced to functional specifications that a

machine can claim to possess. The same report acknowledges elsewhere that the product cannot achieve genuine empathy with people, but this concession appears in a subordinate technical clause in a separate section, while the equivalence claim occupies the main promotional frame. The coexistence of these two statements without reconciliation exemplifies the selective disclosure strategy documented in the greenwashing literature, in which positive claims are foregrounded while limiting information is suppressed (Walker and Wan, 2012; Lyon and Maxwell, 2011). The Mengyu Intelligence Ropet entrepreneurial narrative (T03) produces the equivalence claim through a biological presence mechanism:
"When you stroke its fluffy body, it gazes at you with pure clarity, and its temperature of around 37 degrees makes you feel for a moment that this is a real living creature (T03)."

This formulation constructs AI companion products as functionally equivalent to living beings through the deliberate simulation of biological presence. By engineering body temperature, tactile response, and gaze behaviour to replicate the sensory experience of contact with a living animal, the product asserts that the qualitative experience of companionship can be technically reproduced. The rhetorical substitution conflates the simulated indicators of life with the relational substance of genuine living companionship without making this substitution explicit. Chatzidakis and Littler (2022) identify precisely this move as central to care-washing: the commodification of care converts an ethical relationship into a market transaction, obscuring what is lost in the conversion.

Primary Pattern 2: Vulnerability Erasure.
The second primary pattern, present in nine of twelve texts, involves the systematic omission of elderly vulnerability from product framing. Enterprise texts consistently construct elderly users as capable, digitally empowered, and eager participants in an AI-enabled future rather than as a population whose specific cognitive, social, and economic conditions warrant protective attention. This pattern extends the selective disclosure logic identified in greenwashing research (Marquis et al., 2016) in a specific direction: rather than concealing negative product attributes, enterprises conceal the vulnerability conditions of the target population itself. The Alibaba report (T09) presents the clearest instance:
“Elderly groups are highly capable participants in the AI era and will fully harvest AI dividends across scenarios including home care, health management, nursing supervision, emotional companionship, and career continuation (T09).”

The phrase harvesting AI dividends imports a framework of financial empowerment, positioning elderly people as investors in a technological future. The qualifier fully signals comprehensive benefit capture with no acknowledgement of differential access by age, health status, digital literacy, or income. This framing directly contradicts what Baker et al. (2005) identify as the defining condition of consumer vulnerability: the convergence of cognitive, social, and economic conditions that limit individuals’ capacity to protect their own interests. The promotional frame of universal empowerment makes no allowance for the well-documented relationship between advancing age and declining digital literacy that Czaja et al. (2006) identify, nor for the heterogeneity of the elderly population across cognitive capacity, physical condition, and economic resource. The construction of a uniformly empowered elderly consumer thus functions to suppress precisely the vulnerability conditions that make misleading promotional claims most consequential for this population.

The Mengyu Intelligence Ropet entrepreneurial narrative (T03) constructs the empowered elderly subject through a different mechanism, specifying intact cognition as a product participation criterion. This specification, presented as a design feature ensuring informed consent, simultaneously reveals the product's effective market boundary: elderly users with cognitive impairment, a substantial and disproportionately lonely subgroup, are excluded from the target population. The discourse of empowerment thus applies selectively to the least vulnerable segment of the elderly population while the product is positioned in broader market discourse as a solution to ageing-society care deficits. Moschis (2011) notes that the elderly population is internally heterogeneous in ways that market discourse routinely flattens; the Meela case illustrates how vulnerability erasure operates not through outright falsification but through the selective definition of who counts as the target user.

Structural Condition: Commercial Logic Embedded Within Care Framing.
Both primary patterns are enabled by a structural discursive condition in which commercial and cost-efficiency rationale is embedded within, rather than displaced by, care-oriented language. This condition is what decoupling theory (Meyer and Rowan, 1977) predicts for organisations facing simultaneous commercial and moral legitimacy pressures: formal public discourse conforms to the care norms that the market demands, while the actual commercial logic is maintained in registers addressed to different audiences. The Shounanshan enterprise self-introduction (T01), produced for an industry expo audience, is analytically instructive because it presents this structure with unusual transparency:

"We define our business as helping silver economy enterprises with client relationship maintenance, embedding hardware functions into cooperating enterprises' operations and adding to them a digital employee who can accompany clients at all times. The effect is like how enterprises use DingTalk to manage employees: use Shounanshan to manage clients (T01)."

Addressed to an industry partner audience, the product is presented as a business-to-business client retention tool. The elderly end-user appears as the object of client management rather than the beneficiary of care. The care framing is largely absent. This transparency in an intra-industry register provides evidence that when the intended audience shifts from the general public to commercial partners, the care frame recedes and the commercial logic becomes primary. The contrast between this text and the public-facing texts in the corpus illustrates the audience-dependent nature of the discursive architecture identified as care-washing, and directly confirms the decoupling pattern that Cho and Patten (2007) document in corporate sustainability communication: different registers are deployed for different audiences as a routine feature of organisational legitimacy management, not as exceptional deception.

Table 2
Distribution of Discourse Patterns Across Corpus

Text	Emotional Equivalence	Vulnerability Erasure	Commercial Logic in Care Frame
T01	No	Yes	Yes
T02	Yes	Yes	No
T03	Yes	Yes	Yes
T04	Yes	Yes	No
T05	No	Yes	Yes
T06	Yes	Yes	No
T07	Yes	No	No
T08	Yes	Yes	No
T09	Yes	Yes	No
T10	No	No	Yes
T11	Yes	No	No
T12	No	Yes	No
Total	8/12	9/12	5/12

Note. Codes are non-exclusive. Moral alignment (8/12) and selective disclosure (4/12) are noted as secondary contextual patterns and not reported in full.

Study 2: Thematic Analysis of Practitioner Interviews

Design and Purpose

Study 2 was designed as an explanatory complement to Study 1. Whereas Study 1 identified the discursive structure of AI-mediated care-washing in enterprise promotional communication, Study 2 examines the practitioner-level cognitions that produce and sustain that structure. Specifically, it investigates whether practitioners acknowledge the gap between promotional claims and technical capability, whether they maintain different communicative registers for internal and external audiences, and how they account for the relationship between care discourse and commercial logic. These questions address the mechanism of institutional decoupling from the perspective of those who enact it, rather than from the perspective of the discourse they produce.

Participants and Procedure

Eight practitioners were recruited through purposive sampling targeting individuals with direct operational involvement in AI companion or smart ageing product development and deployment in mainland China and Hong Kong. Participants were identified through the authors' professional networks in the sector. Eligibility required current or recent employment in product planning, operations, market development, or client-facing roles within organisations developing or deploying AI products for elderly users. To ensure the interview data could speak to the discourse structure identified in Study 1 from multiple positions within the production chain, recruitment deliberately spanned upstream product design, downstream market-facing, and cross-functional roles. Table 3 summarises participant profiles.

Table 3
Participant Profiles

Code	Role	Context	Duration
P1	Business development and project operations	AI elderly care technology, mainland China	61 min
P2	Product and business development	Smart ageing and nursing robotics, mainland China	48 min
P3	Product planning and business development	AI elderly care, mainland China	52 min
P4	Smart ageing product operations and marketing	Intelligent elderly care technology, mainland China	47 min
P5	Product planning and design	Smart ageing product design, mainland China	41 min
P6	Project operations and client services	Smart ageing implementation, mainland China	69 min
P7	Product planning and market research	Silver economy market analysis, mainland China	71 min
P8	Residential care home operations	Elderly care home management, Hong Kong	58 min

Note. All participants provided verbal informed consent prior to interview. Interviews were conducted in Mandarin Chinese (P1-P7) and Cantonese-Mandarin (P8) via audio-recorded video call and transcribed verbatim.

Thematic analysis was conducted following Braun and Clarke (2006). An initial deductive framework was derived from the three discourse patterns identified in Study 1, treating those patterns as propositions to be tested against practitioner accounts. Inductive coding was conducted in parallel to capture meanings not anticipated by the a priori framework. Two coders independently coded all eight transcripts, and discrepancies were resolved through discussion until consensus was reached. Themes were defined at the semantic level and reviewed across all eight transcripts for coherence and prevalence. In line with the study's explanatory purpose, recruitment prioritized professional relevance over volume. Thematic saturation was achieved as successive deep-dive accounts yielded conceptual replication without generating new analytical codes (Guest et al., 2006). Three themes met the criteria for full reporting. The coding evolution process, including initial codes, sub-themes, and final themes, is presented in Appendix 1.

Findings

The three themes reported below are organised to correspond directly with the three discourse patterns identified in Study 1. Theme 1 addresses the practitioner-level basis of the emotional equivalence claim. Theme 2 addresses the communicative mechanism sustaining the commercial logic within care framing. Theme 3 addresses the practitioner cognition that corresponds to vulnerability erasure in enterprise discourse.

Theme 1: Acknowledged Technological Misrepresentation

Corresponding to Study 1 Pattern 1 (emotional equivalence claims), Theme 1 captures practitioners' consistent acknowledgement that promotional claims about AI companionship capacity exceed current technical capability. This pattern was present across all eight interviews and constitutes direct practitioner-level evidence that the emotional equivalence asserted in enterprise discourse is not grounded in demonstrated product performance. The convergence of accounts across all eight participants confirms what Abdi et al. (2018) found in their systematic review of AI companion products: that while some products demonstrate benefits in specific functional domains, evidence for sustained emotional and relational outcomes remains limited and methodologically contested. P2 offered the most detailed account, grounded in hands-on care knowledge:

"A lot of promotional materials make people think that robots can already look after elderly people like real people or even replace care workers. But people who have really worked in aged care know that elderly care is far from that simple. (P2)"

This account identifies the source of the gap not as a specifiable missing feature but as the irreducibly situated and relational quality of skilled human care. The gap between what is promoted and what is delivered is characterised as categorical rather than incremental. This is consistent with Vallor's (2011) argument that the moral qualities constitutive of care relationships, including genuine attentiveness, empathy, and responsiveness, cannot be reduced to functional specifications and therefore cannot be replicated by AI systems regardless of technical sophistication. P6, whose role involves direct management of user feedback, provided evidence that the gap also manifests at the point of product use:

"After actually using the AI products, they found that what elderly people really need most of the time is not a device but companionship. Some elderly people liked chatting with the robot, but their interest dropped after a few days. By contrast, when community workers visited them regularly, they were much happier. (P6)"

This observation confirms that the emotional equivalence claim does not survive contact with user experience. Elderly users in P6's account distinguished between device-mediated interaction and human presence and found the former qualitatively insufficient, even after initial engagement. This finding aligns with Sharkey and Sharkey's (2012) concern that the substitution of AI interaction for human relationships in eldercare risks diminishing the relational quality of care even where it improves operational efficiency. P8, operating in Hong Kong's residential care sector, provided an account from a different institutional context that reached the same conclusion:

"Those products generally promote emotional care. But the actual effect, does it really work? The residents here basically just buy a doll or small figurine and hug it as they sit there. (P8)"

The substitution of low-cost comfort objects for marketed AI companion products in P8's account raises a question directly relevant to Study 1's emotional equivalence pattern: whether the emotional needs that AI companion products claim to address are better served through far simpler, non-technological means. This finding does not appear in any of the enterprise texts analysed in Study 1, precisely because it would undermine the promotional logic those texts perform. The absence of such accounts from enterprise discourse is itself evidence of the selective disclosure dynamic that Walker and Wan (2012) identify as central

to legitimacy-seeking communication: information that would undermine the primary legitimacy claim is systematically excluded from the promotional register.

Theme 2: Communicative Register Bifurcation

Corresponding to Study 1's structural condition of commercial logic embedded within care framing, Theme 2 captures practitioners' systematic maintenance of different communicative registers for internal and external audiences. This theme was present in seven of eight interviews and provides the most direct practitioner-level evidence of the institutional decoupling mechanism that produces care-washing as a structural phenomenon rather than individual deception. The pattern confirms what Meyer and Rowan (1977) predict for organisations operating under simultaneous institutional pressures: that formal public accounts will conform to prevailing social norms while actual operational logic is insulated from external evaluation. P1 described the bifurcation in terms of audience-specific emphasis:

"When communicating with clients and users, we talk more about the actual value the product can bring. When communicating internally or with investors, we probably talk more about market scale, industry development trends, future growth in demand due to population ageing, as well as business models and operational efficiency. (P1)"

The same product is presented through a care register to consumer-facing audiences and through a commercial register to investor audiences. This audience-dependent switching is precisely what Cho and Patten (2007) document in corporate sustainability communication: different legitimacy registers are deployed as a routine feature of organisational communication practice rather than as exceptional behaviour. P2 extended this account by identifying a further internal register used exclusively within product development teams:

"But internal discussions are more realistic. We focus on what technology can actually do, which scenarios can genuinely be implemented, and which promotional claims actually exceed current capability. (P2)"

The phrase promotional claims that actually exceed current capability indicates that practitioners are not only aware of the technology-promise gap documented in Theme 1 but that they identify specific instances in which their organisation's external communication crosses from aspiration into misrepresentation. This constitutes a practitioner-level articulation of care-washing as a knowing discursive practice. Bromley and Powell (2012) argue that such decoupling is not experienced by organisational members as deception but as the normal management of different audience expectations, a normalisation that the practitioner accounts in this study consistently confirm. P7 provided the most analytically sharp formulation of the register bifurcation by locating its core axis not in the opposition between care discourse and commercial discourse but in a specifically financial question:

"We are more concerned about is a very practical question: who is actually going to pay. Because a very large ageing market does not necessarily mean a very profitable ageing market. These are two completely different questions. (P7)"

This account introduces a dimension absent from Study 1's discourse analysis: the gap between the size of the market as a demographic fact and extractable revenue within that market as a commercial reality. Enterprise discourse constructs the silver economy as a social problem whose scale justifies AI companionship as a solution; P7's account reveals that within the organisation, the same scale is evaluated primarily as a question of monetisable demand, a calculation that the external discourse systematically conceals. This finding extends Boiral's

(2013) account of sustainability disclosure bifurcation: the gap is not only between social commitment and financial performance, but between the demographic framing of social need and the commercial assessment of market extractability.

Theme 3: Structural Complicity in Vulnerability Concealment

Corresponding to Study 1 Pattern 2 (vulnerability erasure), Theme 3 captures the practitioner-level cognition that the enterprise discourse in Study 1 systematically suppresses. Where enterprise texts construct elderly users as empowered digital consumers, practitioners demonstrate awareness of the specific vulnerabilities that make elderly people susceptible to misleading care claims and express concern about the inadequacy of current industry norms in addressing those vulnerabilities. This awareness corresponds directly to what Baker et al. (2005) identify as the defining condition of consumer vulnerability: cognitive, social, and economic conditions that limit individuals' capacity to protect their own interests in market interactions. P6 provided the most specific account of how vulnerability is exploited through care-oriented promotional discourse:

"Elderly people or their family members spend quite a lot of money to buy them and in the end find that the effect is very different from the promotion. I think this is unfair to consumers. Because many purchasing decisions are actually backed by the filial piety of children. (P6)"

The identification of filial obligation as the mechanism through which promotional discourse converts elderly vulnerability into commercial opportunity is analytically significant. Adult children motivated by care obligations and absence-related guilt are the effective purchasing decision-makers, and they are susceptible to care-oriented promotional language because it aligns with their own emotional motivation. Study 1's vulnerability erasure pattern removes elderly vulnerability from the promotional frame; P6's account reveals that this erasure is commercially functional, because it is precisely the family's emotional response to that vulnerability that drives purchasing. This triangulated dynamic, in which enterprises leverage family emotional motivation rather than engaging directly with elderly users, extends the vulnerable consumer framework of Baker et al. (2005) into a specifically Chinese filial context documented in the research context section (Cong and Silverstein, 2008). P1 identified the structural source of the problem at the level of industry communication norms:

"Sometimes users have very high expectations of AI, but technology actually cannot yet reach that level. So enterprises should try to be objective in their promotions. (P1)"

The normative character of this statement, enterprises should be objective, indicates that P1 does not regard current promotional practice as meeting that standard. This is a practitioner-level critique of care-washing from within the industry, produced in the context of a research interview rather than public discourse, and therefore unavailable through the discourse analysis methods used in Study 1. The gap between P1's private normative standard and the promotional practices his organisation produces is itself a manifestation of the decoupling structure that Bromley and Powell (2012) describe: formal discourse and private cognition are systematically separated as a structural condition of organisational life. P5 articulated a concern that connects vulnerability to the digital inclusion dimension of the empowerment narrative in enterprise discourse:

"Not all elderly people can adapt to digital products. Some elderly people have declining learning capacity or particular physical conditions. If all services become digital, this might actually exclude them. (P5)"

This observation directly counters the vulnerability erasure pattern in Study 1, specifically the Alibaba report's claim that elderly groups will fully harvest AI dividends. P5's account acknowledges the heterogeneity of the elderly population and the risk that technological solutions may systematically exclude the most cognitively and physically vulnerable subgroups, precisely those whose loneliness and care needs are most acute. This concern maps directly onto what Czaja et al. (2006) document regarding the relationship between age, cognitive capacity, and digital literacy: the elderly users most likely to be excluded from AI companion products are precisely those for whom the products' care claims are most socially significant. Moschis (2011) similarly argues that the internal heterogeneity of the elderly population is routinely flattened in market discourse; P5's account demonstrates that practitioners are aware of this flattening even as they participate in producing it.

General Discussion

This study set out to examine how enterprises in China's AI companion product sector publicly construct their products as legitimate care solutions, and whether practitioners who produce that discourse privately understand its relationship to commercial and technical realities. The findings of Study 1 and Study 2 converge on a coherent answer to both questions, and in doing so extend and complicate existing theoretical frameworks in ways that warrant systematic discussion.

Study 1 identified two primary discourse patterns through which AI-mediated care-washing operates. The first, emotional equivalence claims, activates what Suchman (1995) distinguishes as three simultaneous forms of legitimacy: pragmatic legitimacy through assertions of product utility, moral legitimacy through the framing of commercial activity as a response to elderly loneliness, and cognitive legitimacy through the construction of AI companionship as the natural solution to a structural care deficit. The second, vulnerability erasure, extends the greenwashing literature in a direction Lyon and Maxwell (2011) and Delmas and Burbano (2011) did not anticipate: rather than suppressing negative product information, enterprises suppress the vulnerability of the target population itself, replacing the cognitive decline documented by Czaja et al. (2006), the social isolation identified by Alves and Wilson (2008), and the economic constraint noted by Moschis (2003) with a discourse of technological empowerment. Both patterns are enabled by a structural condition in which commercial logic is embedded within care framing, a condition that DiMaggio and Powell (1983) predict as the outcome of field-level mimetic and normative pressures producing homogeneity in public discourse across a sector. Study 2 confirmed and explained these patterns from the practitioner side. Practitioners universally acknowledged a gap between promotional claims and technical capability, attributing it to the irreducibly relational character of human care that Vallor (2011) argues AI systems cannot replicate and that Sharkey and Sharkey (2012) identify as the ethical crux of AI eldercare substitution. They maintained systematically different communicative registers for internal and external audiences, experiencing this bifurcation not as deception but as normal audience management, which is precisely what Palazzo and Scherer (2006) describe as corporate legitimacy communication and what Bromley and Powell (2012) predict as the structural outcome of contradictory institutional demands. And they were privately aware of the elderly vulnerability that enterprise discourse suppresses, including the filial purchasing dynamic in which adult children's care anxiety is converted into purchasing behaviour, a mechanism that

extends Baker et al.'s (2005) vulnerable consumer framework into the Chinese filial context documented by Cong and Silverstein (2008).

These findings collectively support the study's central theoretical contribution. Chatzidakis and Littler (2022) established care-washing as the commodification of care through corporate branding, identifying the first layer of decoupling in which commercial logic is concealed behind care language. The present study demonstrates that AI-mediated care-washing involves a second layer absent from their framework: the concealment of technical uncertainty behind capability claims. Unlike the conventional consumer goods Chatzidakis and Littler examined, AI companion products assert emotional and relational outcomes whose technical basis is empirically contested, as Abdi et al. (2018) document, and whose verification is structurally precluded by the regulatory vacuum identified in the research context. The dual-layer structure, commercial concealment compounded by epistemic concealment, constitutes AI-mediated care-washing as a theoretically distinct phenomenon within the broader taxonomy of corporate washing practices, one whose ethical stakes are amplified by the specific vulnerability of the elderly population it targets and by the filial cultural context that renders that vulnerability commercially actionable.

Beyond the Chinese context, the dual-layer decoupling structure offers a portable analytical lens for other settings in which AI products are marketed as substitutes for relational or affective labour, including AI mental health chatbots and companionship products for isolated populations in other ageing societies. The filial purchasing dynamic is culturally particular, but the underlying mechanism, in which commercial concealment is compounded by epistemic concealment under a regulatory vacuum, is likely to recur wherever AI affective claims outpace verification infrastructure. For policymakers, this indicates the need to treat AI affective claims as a distinct regulatory category, since the harm mechanism operates through discursive substitution rather than product malfunction. For AI developers, limitation acknowledgement should be embedded within product interfaces rather than confined to internal discourse, so that the register bifurcation practitioners currently manage privately becomes a design requirement. For regulatory authorities protecting vulnerable consumers, filial purchasing constitutes a specific site of exposure requiring disclosure standards that address both end user and purchasing decision-maker.

Limitations and Future Research

This study has several limitations that delimit the scope of inference. First, the Study 1 corpus comprises twelve purposively sampled enterprise-affiliated texts, which means it is not exhaustive and may have amplified care-washing patterns by underrepresenting critically oriented or consumer-generated discourse. The Study 2 sample of eight practitioners, while sufficient for thematic saturation (Guest et al., 2006; Morse, 1995), is skewed toward product planning and operations roles in mainland China and cannot support claims about the full population of industry practitioners. Second, both studies examine AI-mediated care-washing exclusively from the production side, through enterprise discourse and practitioner cognition, and do not include data from the elderly consumers who are the target of the discourse being analysed. Baker et al. (2005) argue that understanding consumer vulnerability requires attention to both the structural conditions that create vulnerability and the subjective experiences of vulnerable consumers themselves; the present study addresses only the former. Third, the AI-mediated care-washing construct was developed and tested in the

specific context of China's silver economy, and its generalisability to other national contexts or to other AI application domains making care-adjacent claims, such as mental health AI, AI tutoring products, and grief counselling tools, remains to be established.

Future research should pursue three directions. The corpus-based approach of Study 1 could be extended through systematic sampling of a broader range of text types, including consumer reviews, social media discourse, and regulatory documents, and complemented by computational text analysis methods applied to larger datasets (Blei et al., 2003; Biber and Conrad, 2009). The interview-based approach of Study 2 could be expanded to include practitioners in regulatory, investment, and policy roles. Comparative research across national contexts is a particularly promising direction for establishing the boundary conditions of the dual-layer decoupling structure, including comparative samples from other Asian markets where AI companion products are being deployed in ageing contexts, such as Japan and South Korea (Broadbent et al., 2009; Shibata and Wada, 2011), as well as Western contexts with contrasting regulatory environments and family care norms. Most importantly, future research should incorporate elderly users and their family members as research participants to examine how care-washing discourse is received and acted upon by its intended audience, and whether the vulnerability erasure documented in enterprise discourse corresponds to actual misrecognition of risk on the part of elderly consumers. As AI systems become more capable and their promotional discourse more difficult to evaluate, developing precise criteria for distinguishing legitimate capability claims from care-washing will become increasingly consequential for consumer protection and regulatory design (Floridi et al., 2018; Mittelstadt et al., 2016).

References

Abdi, S., de Witte, L., & Hawley, M. (2020). Emerging Technologies with Potential Care and Support Applications for Older People: Review of Gray Literature. *JMIR aging*, *3*(2), e17286.

Alves, L. M., and Wilson, S. R. (2008). The effects of loneliness on telemarketing fraud vulnerability among older adults. *Journal of Elder Abuse and Neglect*, 20(1), 63-82.

Ashforth, B. E., and Gibbs, B. W. (1990). The double-edge of organizational legitimation. *Organization Science,* 1(2), 177-194.

Baker, S. M., Gentry, J. W., and Rittenburg, T. L. (2005). Building understanding of the domain of consumer vulnerability. *Journal of Macromarketing*, 25(2), 128-139.

Biber, D., and Conrad, S. (2009). *Register, genre, and style*. Cambridge University Press.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993-1022.

Boiral, O. (2013). Sustainability reports as simulacra? A counter-account of A and A plus GRI reports. *Accounting, Auditing and Accountability Journal*, 26(7), 1036-1071.

Braun, V., and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101.

Broadbent, E., Stafford, R., and MacDonald, B. (2009). Acceptance of healthcare robots for the older population: Review and future directions. *International Journal of Social Robotics*, 1(4), 319-330.

Bromley, P., and Powell, W. W. (2012). From smoke and mirrors to walking the talk: Decoupling in the contemporary world. *Academy of Management Annals*, 6(1), 483-530.

Cesta, A., Cortellessa, G., Orlandini, A., and Tiberio, L. (2016). Long-term evaluation of a telepresence robot for the elderly: Methodology and outcomes. *International Journal of Social Robotics*, 8(3), 421-441.

Chatzidakis, A., and Littler, J. (2022). An anatomy of carewashing: Corporate branding and the commodification of care during Covid-19. *International Journal of Cultural Studies,* 25(3-4), 268-286.

Cho, C. H., and Patten, D. M. (2007). The role of environmental disclosures as tools of legitimacy: A research note. *Accounting, organizations and society*, 32(7-8), 639-647.

Cho, C. H., Guidry, R. P., Hageman, A. M., and Patten, D. M. (2012). Do actions speak louder than words? An empirical investigation of corporate environmental reputation. *Accounting, Organizations and Society*, 37(1), 14-25.

Cong, Z., and Silverstein, M. (2008). Intergenerational support and depression among elders in rural China: Do daughters-in-law matter? *Journal of Marriage and Family*, 70(3), 599-612.

Czaja, S. J., Charness, N., Fisk, A. D., Hertzog, C., Nair, S. N., Rogers, W. A., and Sharit, J. (2006). Factors predicting the use of technology: Findings from the Center for Research and Education on Aging and Technology Enhancement (CREATE). *Psychology and Aging*, 21(2), 333-352.

Delmas, M. A., and Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, 54(1), 64-87.

DiMaggio, P. J., and Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147-160.

Drolet, A., Schwarz, N., and Yoon, C. (2019). The influence of aging on consumer decision-making. *Consumer Psychology Review*, 2(1), 3-16.

Fereday, J., and Muir-Cochrane, E. (2006). Demonstrating rigor using thematic analysis: A hybrid approach of inductive and deductive coding and theme development. *International Journal of Qualitative Methods*, 5(1), 80-92.

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., and Vayena, E. (2018). An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689-707.

Guest, G., Bunce, A., and Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59-82.

Guo, M., Chi, I., and Silverstein, M. (2019). The structure of intergenerational relations in rural China: A typology of supporting elderly parents. *Journal of Marriage and Family*, 81(5), 1111-1125.

Hanleybrown, F., Kania, J., and Kramer, M. (2012). Channeling change: Making collective impact work. *Stanford Social Innovation Review*, 10(1), 1-8.

Ivankova, N. V., Creswell, J. W., and Stick, S. L. (2006). Using mixed-methods sequential explanatory design: From theory to practice. *Field Methods*, 18(1), 3-20.

Laufer, W. S. (2003). Social accountability and corporate greenwashing. *Journal of Business Ethics*, 43(3), 253-261.

Lyon, T. P., and Maxwell, J. W. (2011). Greenwash: Corporate environmental disclosure under threat of audit. *Journal of Economics and Management Strategy*, 20(1), 3-41.

Marquis, C., Toffel, M. W., and Zhou, Y. (2016). Scrutiny, norms, and selective disclosure: A global study of greenwashing. *Organization Science*, 27(2), 483-504.

Meyer, J. W., and Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology*, 83(2), 340-363.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data and Society*, 3(2), 1-21.

Moschis, G. P. (2003). Marketing to older adults: An updated overview of present knowledge and practice. *Journal of Consumer Marketing*, 20(6), 516-525.

Moschis, G. P. (2011). Research frontiers on older consumers' vulnerability. *Journal of Consumer Affairs*, 45(3), 467-491.

National Health Commission of China. (2022). *Guiding opinions on promoting the development of intelligent health and elderly care services*. National Health Commission.

Palazzo, G., and Scherer, A. G. (2006). Corporate legitimacy as deliberation: A communicative framework. *Journal of Business Ethics*, 66(1), 71-88.

Ramus, C. A., and Montiel, I. (2005). When Are Corporate Environmental Policies a Form of Greenwashing? *Business & Society*, 44(4), 377-414.

Roberts, H., Cowls, J., Morley, J., Taddeo, M., Wang, V., and Floridi, L. (2021). The Chinese approach to artificial intelligence: An analysis of policy, ethics, and regulation. *AI and Society*, 36(1), 59-77.

Sharkey, A., and Sharkey, N. (2012). Granny and the robots: Ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1), 27-40.

Shibata, T., and Wada, K. (2011). Robot therapy: A new approach for mental healthcare of the elderly. *Gerontechnology*, 10(3), 140-147.

Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review*, 20(3), 571-610.

Suta, P., Chan, J. H., and Mongkolnam, P. (2025). Enhancing Elderly Well-Being Through AI-Based Virtual Companion Systems: A Cooperative Approach. *In International Conference on Cooperative Design, Visualization and Engineering (pp. 86-99)*. Cham: Springer Nature Switzerland.

Vallor, S. (2011). Carebots and caregivers: Sustaining the ethical ideal of care in the twenty-first century. *Philosophy and Technology*, 24(3), 251-268.

Walker, K., and Wan, F. (2012). The harm of symbolic actions and green-washing: Corporate actions and communications on environmental performance and their financial implications. *Journal of Business Ethics*, 109(2), 227-242.

Zhan, H. J., and Montgomery, R. J. V. (2003). Gender and elder care in China: The influence of filial piety and structural constraints. *Gender and Society*, 17(2), 209-229.

Zhan, H. J., Feng, Z., and Luo, B. (2015). Recent developments in institutional elder care in China: Changing concepts and attitudes. *Journal of Aging and Social Policy*, 20(2), 126-148.