

MASTER'S THESIS

Reevaluating Technology Ideologies:

How Consumer Narratives about Humanoid Robots Challenge or Affirm Kozinets' Ideological Field Model in an Accelerating Tech Era

submitted by

BSc (WU) Daniel Eric Novas Alonso

intended degree

Master of Science (WU), MSc (WU)

Student ID number: 11924600

Degree programme: Master's Programme in Marketing

Supervisor: Prof. Dr. Marius Lüdicke

Vienna, June 2026

Abstract

This thesis examines how contemporary consumer narratives about humanoid robots align with, challenge, or extend Kozinets' (2008) ideological field model of technology consumption. Drawing on twenty-four semi-structured in-depth interviews conducted between February and March 2026 with respondents of diverse professional, generational, and national backgrounds, the study employs a hermeneutic part-to-whole interpretive procedure that combines deductive engagement with the original four-element framework (Techtopian, Green Luddite, Work Machine, and Techspressive) and inductive openness to emergent patterns. Six cross-cutting patterns structure the main analysis: a qualified optimism paradox, surveillance as a master fear, geopolitical and corporate provenance as a structural trust variable, velocity anxiety, human-made value as a resistance ideology, and cognitive atrophy as an emergent concern. A further six patterns, among them the home as the sharpest ideological cleavage, the work-purpose distinction, and the robot as a projective social mirror, are documented in an appendix. The original four-element model remains valuable, but five structural extensions are proposed: a fifth ideological element organised around the irreducibility of the human, two substrates (surveillance and provenance) that condition all elements, a velocity dimension, and the relocation of cognitive atrophy as a cross-cutting inner contradiction. This thesis closes with a phenomenological coda reading the post-interview reversal through Gabriel Marcel's philosophical structure of intersubjective hope.

Keywords: humanoid robots, technology ideology, consumer narratives, Kozinets, ideological field model, Gabriel Marcel, hope

List of Figures

Figure 1 Kozinets' (2008) ideological field model.....	9
Figure 2 The revised ideological field model introduced in three layers. Elements and semantic relations adapted from Kozinets (2008).....	23
Figure 3 The orbital ideological field model of technology consumption in an accelerating tech era. Elements and semantic relations adapted from Kozinets (2008).	52

List of Tables

Table 1 Respondent composition	16
Table 2 The five elements of the revised model along identical dimensions.	54

Table of Contents

Abstract.....	ii
List of Figures.....	iii
List of Tables.....	iv
1 Introduction	1
2 Literature Review	4
2.1 Technology, Consumption and Ideology.....	4
2.2 Kozinets' (2008) Ideological Field Model	6
2.2.1 Theoretical Antecedents	6
2.2.2 The Four Ideological Elements	6
2.2.3 Semantic Relations and Consumer Dynamics.....	8
2.3 Related Frameworks and Extensions.....	10
2.4 Consumer Research on Humanoid and Social Robots.....	10
2.5 The Research Gap and Theoretical Provocation	11
3 Method.....	13
3.1 Research Paradigm and Methodological Orientation	13
3.2 Research Design	14
3.3 Sampling Strategy and Respondent Composition.....	14
3.4 Data Collection.....	18
3.5 Ethical Considerations	19
3.6 Researcher Positionality and Reflexivity.....	20
3.7 Data Analysis	20
3.8 Methodological Limitations	22
4.1 The Qualified Optimism Paradox	24
4.1.1 Coping Strategies and the Architecture of Optimism	26
4.2 Surveillance as Master Fear	27
4.3 Geopolitical and Corporate Origin as Trust Proxy	29
4.4 Velocity Anxiety: Pace as Ideological Provocation.....	31

4.5 Human-Made Value as Resistance Ideology	32
4.6 Cognitive Atrophy as Emergent Anxiety	34
4.7 Synthesis: Configurations and Cleavages	36
5 Discussion	38
5.1 What the Data Affirm	38
5.2 What the Data Challenge	40
5.3 What the Data Extend	41
5.3.1 The Surveillance Substrate	41
5.3.2 The Provenance Variable	42
5.3.3 The Velocity Dimension.....	44
5.3.4 Cognitive Atrophy as a Fifth Inner Contradiction	46
5.3.5 The Human-Reservation Position.....	48
5.4 A Revised Model	51
5.5 Theoretical Implications	55
6 Conclusion.....	58
6.1 Summary of Contributions.....	58
6.2 Practical Implications	59
6.3 Limitations Revisited	62
6.4 Directions for Further Research	63
6.5 Closing Reflection.....	65
6.6 A Theoretical Coda: The Interview as Hope-Event.....	66
References.....	70
Appendix A: Interview Guide	76
A.1 Pre-Interview Procedure	76
A.2 Section 1: Warm-Up and Introduction (5–7 minutes).....	76
A.3 Section 2: Open Associations and Feelings (10–12 minutes)	77
A.4 Section 3: Personal Expectations and Sentiments (12–15 minutes)	77
A.4.1 Optional Tailored Probes	77
A.5 Section 4: Deeper Reflections and Mixed Feelings (8–10 minutes)	78

A.6 Section 5: Information Sources and Closing (7–10 minutes)	78
A.7 Closing Procedure	79
<i>Appendix B: Use of Generative Artificial Intelligence (AI)</i>	<i>80</i>
<i>Appendix C: List of Aids Used.....</i>	<i>81</i>
<i>Appendix D: Plagiarism</i>	<i>82</i>
<i>Appendix E: Respondent Profiles</i>	<i>83</i>
E.1 Expert Respondents.....	83
E.2 Consumer Respondents	85
<i>Appendix F: Additional Narrative Patterns.....</i>	<i>87</i>
F.1 The Domestic Sphere as Ideological Boundary	88
F.2 The Uncanny Valley as Ideological Trigger	89
F.3 The Work-Purpose Distinction and the UBI Cleavage	91
F.4 The Emotional Attachment Split.....	94
F.5 The Generational Gradient	97
F.6 The Robot as Social Mirror	99
F.6.1 The Interview as Symptom: A Search for Authority	101

1 Introduction

In early 2026, footage of humanoid robots performing synchronised martial arts routines during China's Spring Festival Gala circulated widely across television and social media platforms (Al Jazeera, 2026; CNN, 2026; Reuters, 2026). The robots executed backflips and interacted with human performers with a fluidity that, just twelve months prior, would have seemed years away. For the twenty-five participants of this study, many of whom had encountered precisely this footage in the days before their interview, the clip was a visceral marker of technological acceleration. Machines that had struggled to walk in a straight line were now performing backflips on a public stage, in a foreign country. More than once, respondents described the same arc of surprise. One participant, working in Capital Markets, captured the collective reaction: "what surprised us the most is the pace of progress."

This moment captures the central challenge this thesis addresses. Humanoid robots, autonomous bipedal machines designed to perceive, navigate, and act within environments built for human bodies, are no longer bound to science fiction or academic robotics laboratories. Platforms such as Tesla's Optimus, Boston Dynamics' Atlas, Unitree's G1 and Figure's 01 are advancing toward commercial deployment in manufacturing, logistics, and eventually domestic settings. Hence, the question of how consumers think, feel, and narrate their relationship with this technology has moved from speculative to urgent. Consumer researchers, marketers, and policymakers therefore need frameworks to map the ideological terrain through which humans make sense of a technology that might become omnipresent in their working and domestic life within the next decade.

The most comprehensive framework available for this purpose is Kozinets' (2008) ideological field model of technology consumption. Developed through an in-depth qualitative investigation of consumer narratives, the model identifies four ideological elements: Techtopian (technology as societal progress), Green Luddite (technology as infringement on the natural), Work Machine (technology as productivity increase), Techspressive (Technology as self-expressive medium to derive pleasure from).

Kozinets maps the semantic relations between these four elements onto a semiotic square. The framework revealed that consumers do not occupy fixed ideological positions but move fluidly between elements, driven by inner contradictions within each ideology and the contrasts between them. Published in the *Journal of Consumer Research*, it remains a systematic account of how technology ideologies influence how consumers think and express themselves.

Kozinets' model nonetheless predates large language models, the recent acceleration of robotics, and the present environment in which the pace of change has itself become overwhelming for many consumers. The six informants whose narratives formed the empirical basis of the original study were navigating a landscape of desktop computers, early internet adoption, and Palm Pilots. The technologies were relatively mature and the ideological terrain was comparatively settled. Especially the pace of change was perceptibly slower than the conditions under which the consumers in the present study are operating. Eighteen years on, the question therefore arises whether Kozinets' ideological field model retains its analytical validity under the technological and temporal conditions of the contemporary moment.

The question this thesis pursues is straightforward: does Kozinets' framework still adequately capture how consumers narrate their relationship with technology when the technology in question is not just new, but paradigm-shifting, physically embodied, and arriving at a pace that outstrips imagination?

This thesis approaches that question empirically. Over the course of February and March 2026, twenty-four semi-structured in-depth interviews were conducted, comprising twenty-one consumer respondents and four experts. The consumer sample was purposive, drawn to span diverse professional backgrounds, age groups, and national contexts: professionals from Capital Markets, computer forensics, molecular biology, law, marketing, hospitality management, and diplomatic service, alongside academics, entrepreneurs, and students. The expert interviews were conducted with individuals operating at the frontier of this technological era: the founder of Tesla, Owners of Silicon Valley, the CIO of an Austrian industrial construction group actively integrating robotic systems, an AI and marketing professor at WU Vienna, and a doctoral researcher who lives and works alongside a domestic humanoid robot called Tova. These conversations form an ideological dataset that captures a range of beliefs from profound technological enthusiasm to principled scepticism.

The central research question guiding the analysis is: how do contemporary consumer narratives about humanoid robots align with, challenge, or extend Kozinets' (2008) ideological field model of technology consumption? Three subsidiary questions structure the investigation. First, to what extent do the four ideological elements manifest in consumer narratives about humanoid robots, and in what configurations? Second, what patterns of ideological movement, contradiction, and co-articulation emerge across a sample that spans diverse professional groups?

Third, does the data surface emergent patterns, such as temporal anxiety, surveillance concerns, the domestic sphere as an ideological boundary, or the geopolitical origins of technology, that the 2008 model does not adequately capture?

The thesis makes three primary contributions. Theoretically, it constitutes, in the view of the author, the first systematic empirical test of Kozinets' ideological field model in which humanoid robots are the focal technology, and it does so in a period of documented acceleration in the field. Empirically, it documents a set of cross-cutting patterns in contemporary consumer narratives about humanoid robots, surfacing both the model's strengths and the areas in which it may require modification or extension. Methodologically, the thesis contributes a five-cycle hermeneutic interpretive procedure that combines deductive engagement with an established theoretical framework and inductive openness to emergent patterns, and that can be applied to comparable studies of ideological articulation in adjacent consumption domains.

The remainder of the thesis proceeds as follows. Chapter 2 reviews the relevant literature and develops the theoretical structure for the study by examining the concept of technology ideology in consumer research, explicating Kozinets' (2008) model in detail, and situating it alongside related frameworks and subsequent extensions. Chapter 3 describes the methodological approach, including the sampling strategy, interview design, analytical procedure, and the rationale for key methodological decisions. Chapter 4 presents the empirical findings, organised around six cross-cutting patterns identified in the dataset, with a further six recorded in Appendix F. Chapter 5 discusses these findings in relation to the theoretical literature, draws out the implications for the model's adequacy in the present moment, and proposes a revised model. Chapter 6 consolidates the contributions of the study while developing practical implications and acknowledges the limitations. On top of that, directions for further research are being presented. References and appendices, including the interview guide, follow.

2 Literature Review

This chapter develops the theoretical foundation for the study. It begins by examining how ideology has been treated within consumer research, particularly in relation to consumer narratives about technology. It then explains Kozinets' (2008) ideological field model, covering its intellectual precursors, its four ideological elements, the semantic relations between them, and the model's findings on consumer fluidity. The chapter then turns to related frameworks and extensions, including the technology paradoxes identified by Mick and Fournier (1998) and Thompson's (2004) account of marketplace mythologies. This is followed by a brief review of consumer research on humanoid and social robots. The chapter concludes by articulating the research gap this thesis addresses.

2.1 Technology, Consumption and Ideology

The concept of ideology has a long and contested history in social theory, but within consumer research it denotes something specific: the system of meanings, values, and beliefs through which individuals and groups make sense of their material circumstances and social relations, and which shapes how they act upon the world (Thompson, 1990; Belk, Wallendorf, & Sherry, 1989). Ideology is not just a set of explicitly held beliefs; it operates largely through what feels natural, inevitable, or simply "how things are." Thompson (1997) argues in his account of ideology and modern culture that dominant ideological formations sustain relations of power precisely by appearing to be common sense, instead of contested constructions.

Within marketing and consumer culture theory (CCT), ideology has been studied primarily as a force that shapes the meanings consumers attach to products, brands, and consumption practices (Arnould & Thompson, 2005). Of particular relevance are the ideological dimensions of technology consumption. Mick and Fournier (1998) identified eight paradoxes that structure consumer experiences of technologies. Hence, technology consumption is never just a matter of utility but is always filled with meanings, values, and contradictions. Thompson (2004) examined what he termed marketplace mythologies around technology and demonstrated that American consumers tend to inhabit an ambivalent ideological space that combines quasi-religious faith in technological progress with deep anxieties such as: loss of control, authenticity, or human connection.

What was missing from this literature, as Kozinets (2008) noted, was a systematic map not just of individual paradoxes or myths but of the organised structure of ideological elements through which technology consumption is articulated, together

with the relational logic that drives consumers to move between them. This gap motivated his original publication, which forms the theoretical backbone of the present thesis.

Within consumer research specifically, Mick and Fournier's (1998) study of consumer cognizance, emotions, and coping strategies in relation to technology is a direct theoretical bridge between the broader sociology of technology surveyed above and the ideological-field analysis this thesis pursues.

Their account identifies eight paradoxes that technology consumption activates: Control/chaos, Freedom/enslavement, New/obsolete, Competence/incompetence, Efficiency/inefficiency, Fulfills/creates needs, Assimilation/isolation, Engagement/disengagement.

They also identify the coping strategies through which consumers manage to deal with these paradoxes: pre-purchase intelligence and avoidance before the encounter, then confrontation, partnering, mastery, accommodation, or neglect afterwards. The 1998 framework is a precursor to Kozinets' (2008) ideological field model: both conducted in-depth interviews, both treat consumer narratives as the main site of theoretical work, and both treat contradiction within the consumer as a structuring feature rather than inconvenience of data. A core difference between them is that Mick and Fournier emphasise the paradox-coping cycle as the engine of the consumer-technology encounter, whereas Kozinets focusses on the field of ideological positions through which consumers move. Both frameworks are mutually clarifying rather than competing, and Chapter 4 returns to Mick and Fournier (1998) explicitly to interpret the coping mechanisms surfaced in the corpus.

Three additional theories are relevant. Borgmann's (1984) "device paradigm" explains how modern consumer technology tends to replace traditional tools. While older tools required active, engagement-rich "focal practices" to shape meaning, work, and relationships, modern devices prioritize convenience and separate the user from deeper human experiences. Borgmann's vocabulary is useful to analyse consumer resistance to technological substitution. Chapter 4 refers to it for the human-made-value pattern and the domestic-sphere defence. Verbeek's (2005, 2011) postphenomenological framework, developed within the philosophy of technology lineage of Ihde (1990), complements Borgmann and analyses how technologies *mediate* the relationship between user and the world. Essentially, not as neutral devices but as active configurers of perception and moral positions. This was applied directly to the projective-mirror pattern identified in Chapter 4. Here the humanoid

robot operates as the mediating object through which respondents articulate prior ideological content.

Belk's (1988) account of the extended self provides the consumer-research complement to these philosophical accounts as a useful framework to understand the differential affective load that the humanoid form carries.

2.2 Kozinets' (2008) Ideological Field Model

2.2.1 Theoretical Antecedents

Kozinets' (2008) model is grounded in a combination of critical discourse analysis, articulation theory, and the Greimasian semiotic square. From discourse theory (Laclau & Mouffe, 1985; Žižek, 1989), the model draws the concept of an ideological field constituted as a montage of "floating signifiers". These are chains of meaning whose content is not fixed but is retroactively anchored by nodal points. A nodal point is essentially a semiotic anchor that fixes the meaning of surrounding signifiers relative to itself. Articulation refers to the practice through which ideological elements are connected to nodal points, creating distinctive subject positions from which individuals can interpret experience and act. When a consumer "interpellates" an ideology, a term coined in Althusser (1971), they assume a subject position from which that ideology makes the world intelligible and their place within it secure.

The semiotic square (Greimas, 1966; Jameson, 1981) provides the spatial logic of the model. It maps two axes of opposition to generate a field of four positions. These positions stand in three types of semantic relation to one another: contrariety (opposition), contradiction (negation), and complementarity (implication).

Jameson (2005) emphasised that the semiotic square is particularly well suited to capture ideological formations because it depicts the paradox driving desire for an "impossible synthesis". What he means by this is the consumer's unrealisable longing for a resolution that would reconcile the contradictions within the field. It is precisely this unrealisable desire that fuels the dynamic movement between ideological positions that Kozinets observes in his data.

2.2.2 The Four Ideological Elements

Kozinets identifies four ideological elements that together constitute the ideological field of technology consumption. Each element is organised around a nodal point and

each contains within itself a set of inner contradictions that produce instability and movement.

The Techtopian element articulates technology with the supreme good of Progress. Historically, it draws from the Renaissance utopianism of thinkers such as Sir Thomas More and Francis Bacon, both of whom linked scientific and technological advancement with social and moral betterment. Also the progressive idealism of the 1960s space programme era fuels Techtopian ideology. In its contemporary form, Techtopian ideology is ubiquitous in mainstream business discourse and popular science communication. The emotional register of Techtopian ideology is the technological sublime. An experience of awe, transcendence, and quasi-religious wonder in the face of technological achievement (Nye, 1994; Noble, 1999). Its inner contradictions are significant, however: the overly optimistic and moral tone of Techtopian discourse sits in tension with the historical record of technological disasters (Hiroshima, Chernobyl), the inaccessibility of its promises to those with little resources, and with the cynicism generated by successive waves of technological revolution that have failed to deliver on their utopian promises (Frank, 2000; Markoff, 2005).

The Green Luddite element articulates technology with the destruction of the supreme good of Nature and traditional ways of life. Its historical roots lie, as the name suggests, in the Luddite movement. The Luddites were an organised resistance of English textile workers to early industrial machinery. It has been sustained by a diverse ideological coalition: Amish and Quaker communities, beatniks and hippies, downshifTERS, voluntary simplicity advocates, greens, ecofeminists, and anti-globalisation activists. Interestingly, the Green Luddite ideology offers the only position within the field from which to oppose technology consumption. It is also, as Kozinets notes, stigmatically associated with unfashionable, austere, or violent historical actors. Its inner contradictions include the difficulty of avoiding technology altogether in a technologically mediated world, a lack of viable political purchase for anti-technology positions, and the association with losing groups. The Luddites themselves were on the losing side of history.

The Work Machine element articulates technology with the supreme good of Economic Growth and productivity. Its origins run through the coincident development of capitalism and industrial technology in the eighteenth and nineteenth centuries, through Frederick Winslow Taylor's (1911) scientific management and its articulation of efficiency as the highest organisational virtue, as well as through the dominance of professional life by the computer and internet in the information economy. Work Machine ideology treats technology consumption from an unemotional, instrumental

perspective: technology is valuable to the extent that it increases output, reduces costs, and confers competitive advantage. Its inner contradictions are the shadow side of industrial capitalism. Together with the golden promises of productivity and prosperity, Work Machine ideology resonates with the overtones of enslavement, exploitation, conformity, and loss of control identified by critical theorists from Ellul (1964) to Heidegger (1962/1927).

The Techspressive element articulates technology with the supreme good of Pleasure and self-expression. It is historically the most recent of the four, emerging from the global spread of video game culture since the 1970s and the rise of what Rushkoff (1996) and Tapscott (1998) called the "game generation." In the 1990s, geek chic and the culture of Wired magazine crystallised an ideology in which cutting-edge technology consumption was a form of identity expression, a signal of creativity and coolness, becoming itself a cultural currency (Turner, 2006; Streeter, 2011). Contemporary Techspressive ideology encompasses the aestheticisation of smartphones and tech-devices to the indulgence of gaming, social media, and self-documentation through digital tools. Its inner contradictions lie in the fragility and impermanence of technologically mediated pleasure: the liberation of digital play risks addiction and possible isolation.

2.2.3 Semantic Relations and Consumer Dynamics

The four elements are organised in a semiotic square through four axes of semantic relation. Techtopian and Green Luddite stand in a contrariety of morality: both are collectively oriented, morally serious ideological positions that disagree fundamentally about whether technology consumption is good or bad for society. Techtopian and Techspressive stand in a contrariety of individualism: both endorse technology but disagree about whether its supreme good is social progress (Techtopian) or personal pleasure (Techspressive). Work Machine and Green Luddite stand in a contrariety of standards. Both ideologies acknowledge technology's social effects but disagree about the supreme standard: economic growth or natural ways of life. Work Machine and Techspressive stand in a contrariety of indulgence: both are comparatively individualistic but differ in their relationship to pleasure, the former treating technology instrumentally and the latter treating it as an end in itself. Additionally, Techtopian and Work Machine ideology share a complementarity of reason (both are grounded in Enlightenment and functionalist ideals of rational calculation), while Green Luddites and Techspressives share a complementarity of emotion (both involve passionate, non-calculative relationships to technology, fear and pleasure).

Figure 1 below presents the model visually, reproducing the four-element semiotic square with each element's supreme good and its four inner contradictions.




Figure 1 Kozinets' (2008) ideological field model.

Kozinets demonstrates that these four ideological elements are not fixed positions that consumers occupy permanently, but rather nodes through which they move in response to contextual triggers, inner contradictions, and the competing demands of their life circumstances. The six informants in his original study all exhibit within-interview movement across multiple ideological positions: a salesperson shifts between Work Machine pragmatism and Green Luddite romanticism as her standards of personal fulfilment change; a Silicon Valley programmer moves unsteadily between Techtopian social idealism and Green Luddite concern about surveillance; a web designer oscillates between Work Machine instrumentalism and Techspressive pleasure in a rhythm driven by the moral injunction that he “should” be working rather than playing. The dynamic quality and the model’s account of ideological fluidity is among its most distinctive theoretical contributions.

2.3 Related Frameworks and Extensions

Kozinets' model did not emerge in isolation. Several adjacent frameworks extend the theoretical terrain on which it operates. Mick and Fournier (1998), discussed at length in Section 2.1, focused primarily on the phenomenological level of individual technology encounters rather than the macrosocial level of ideological formation. Their paradox framework anticipates many of the inner contradictions Kozinets later formalises within his model: the tension between Work Machine efficiency and its shadow of enslavement, for instance, maps closely onto their efficiency-inefficiency and freedom-enslavement paradoxes.

Thompson (2004) offered an account of American technology mythology as fundamentally ambivalent: consumers simultaneously inhabit a quasi-Romantic belief in the redemptive power of technological progress (associated with the Techtopian) and a resistance to industrial modernity (associated with the Green Luddite). Thompson characterised this as an "ambivalent and even schizoid" cultural relationship to technology, drawing on the same Techtopian/Green Luddite opposition that forms one axis of Kozinets' model. Kozinets' demonstrated that this binary account was insufficient: The two additional ideological elements (Work Machine and Techspressive) are necessary to account for the full range of consumer articulations. Thus, the model's relational structure explains why consumers so frequently appear contradictory or inconsistent in their technology attitudes.

2.4 Consumer Research on Humanoid and Social Robots

Empirical consumer research on humanoid and social robots has grown substantially since 2015, but it concentrates on a narrow set of questions. A substantial body of work has examined the antecedents and dimensions of consumer acceptance, drawing on the Technology Acceptance Model (TAM) and its extensions (Heerink et al., 2010; de Graaf & Ben Allouch, 2013). Research on the uncanny valley (Mori, 1970; MacDorman & Ishiguro, 2006) has explored the relationship between a robot's degree of human likeness and consumer comfort. People usually feel that robots which closely but imperfectly resemble humans elicit heightened discomfort and dissonance. Studies of service robot adoption have examined consumer responses to robots in hospitality, retail, and healthcare contexts (Wirtz et al., 2018; Ivanov & Webster, 2019); here the acceptance varies significantly by service type, consumer age, and cultural context.

A smaller body of work has examined the emotional and relational dimensions of human-robot interaction. Turkle (2011) offered an influential account of the psychological mechanisms through which people form attachments to robots and AI

systems. He argues that such attachments exploit evolved tendencies toward anthropomorphism and that their widespread adoption could have significant consequences for the texture of human social and emotional life. Research on companion robots has demonstrated that emotional attachments to robotic companions can be meaningful and psychologically beneficial for isolated or elderly populations, while strongly raising ethical questions about the authenticity of such relationships (Sharkey & Sharkey, 2012).

The existing literature on humanoid robot consumption has rarely engaged with the ideological dimensions of consumer narratives in the way Kozinets' framework affords. Studies tend to examine specific attitude constructs such as trust, anthropomorphism, perceived usefulness or anxiety, rather than the broader ideological formations through which consumers make sense of robots as a cultural and political phenomenon. This leaves a significant gap: we know how consumers evaluate specific robot features in particular contexts, but we know relatively little about how they ideologically position themselves in relation to humanoid robots as a class of transformative technology, and especially how those positions relate to one another in a structured field.

2.5 The Research Gap and Theoretical Provocation

This thesis addresses a gap that lies at the intersection of two observations. The first is that Kozinets' (2008) ideological field model was developed in a technological and historical context significantly different from the present moment. His study was conducted in an era of desktop computers and early internet adoption. The mapped ideological field was dynamic but not destabilised and the inner contradictions and semantic relations it identified, were the product of decades of accumulated cultural sedimentation around mature technologies.

The second observation is that humanoid robots are a categorically different kind of technological object: An embodied technology that concentrates an unusual range of ideological provocations in a single form. Unlike a laptop or a smartphone, a humanoid robot operates in intimate spaces, mimics the human form, has the potential to displace human labour at scale, and raises questions about consciousness, rights, and the boundaries of the human that no previous consumer technology has posed so directly. It arrives, moreover, not gradually but with what multiple participants in this study described as a disorienting speed. The pace of development has itself become a

source of anxiety, quite apart from any concern about what the technology does or might do.

Together, these observations motivate a research question that is both an empirical inquiry and a theoretical test: does Kozinets' ideological field model adequately capture the ideological dimensions of consumer narratives about humanoid robots in an era of technological acceleration? Or does the data surface patterns that require the framework to be extended, modified, or supplemented?

The following chapters develop an empirical response to these questions.

3 Method

This chapter sets out the methodological foundations of the empirical study. It begins with the interpretive paradigm that informs the research and the reasons for a qualitative, in-depth interview design. It then describes the sampling strategy and respondent composition, how the interviews were conducted, the procedures for ethical compliance and data protection, and the researcher's positionality. The chapter closes with a detailed account of the analytical procedure, a hermeneutic part-to-whole interpretation that combines deductive engagement with Kozinets' (2008) ideological field model and inductive openness to emergent themes. An evaluation of the study's trustworthiness and limitations follows.

3.1 Research Paradigm and Methodological Orientation

The study sits within an interpretive research paradigm and draws on the hermeneutic tradition of qualitative inquiry that has shaped consumer culture theory (CCT) since the 1990s (Arnould & Thompson, 2005). Its underlying ontological assumption is that consumer narratives about technology are not direct reports of stable inner attitudes. They are socially situated articulations through which respondents actively construct meaning, identity, and moral positions in relation to the world they inhabit (Thompson, 1997). The corresponding epistemological commitment is that such meanings are intersubjectively accessible. They can be reconstructed, with care and reflexivity, through the dialogic encounter of the interview and the interpretive reconstruction of the resulting text (Spiggle, 1994; McCracken, 1988).

The object of inquiry is the structure and dynamics of an ideological field, a configuration of meanings, values, and contradictions, rather than the distribution of variables across a population. The appropriate strategy is therefore one that produces depth, narrative texture, and cross-case interpretability rather than statistical generalisability (Belk, Wallendorf, & Sherry, 1989; Kozinets, 2008). The study adopts a qualitative, interpretive design grounded in semi-structured in-depth interviews. Analysis proceeds through iterative hermeneutic engagement with the resulting transcripts.

Kozinets' (2008) ideological field model was itself developed from a small, purposive sample of in-depth interviews. He analysed them through a hermeneutic procedure informed by critical discourse analysis and the Greimasian semiotic square. To test, challenge, and potentially extend that model in a new technological and historical context, it is methodologically coherent (and analytically conservative) to employ a similar mode of inquiry rather than a categorically different one. The present study

can therefore be understood as a partial replication and theoretical extension of Kozinets' original design. It holds the interpretive paradigm constant while varying the technological focus (humanoid robots in place of personal computing) and the historical moment (the mid-2020s in place of the early 2000s).

3.2 Research Design

The empirical material consists of twenty-four semi-structured in-depth interviews conducted between February and March 2026. The interview was selected as the primary method because the study's research questions concern how consumers articulate, narrate, and reason about a technology that most of them have not yet encountered in everyday life. Survey instruments would have constrained respondents to predetermined response categories. They could not have surfaced the contradictions, hesitations, and qualified positions that are themselves of theoretical interest. Ethnographic observation was not feasible because consumer-grade humanoid robots are not yet present in domestic or workplace settings on a meaningful scale. The in-depth interview, by contrast, allows respondents to speak at length, to revise and contradict themselves, and to articulate moral and emotional positions whose internal structure can subsequently be reconstructed from the transcript (McCracken, 1988; Thompson, Locander, & Pollio, 1989).

The semi-structured format was chosen deliberately over more rigid alternatives. A common interview guide ensured that all respondents addressed a comparable set of thematic areas, supporting cross-case comparison. However, departing from the guide allowed the researcher to follow emotionally or theoretically significant threads as they arose, to probe contradictions, and to invite respondents to develop their own framings rather than merely react to the researcher's prompts. This balance of structure for comparability and openness for depth is consistent with established qualitative practice in CCT (Arnould & Wallendorf, 1994).

3.3 Sampling Strategy and Respondent Composition

Respondents were recruited through a combination of purposive and snowball sampling. Purposive sampling was used to construct a respondent pool that varied systematically across dimensions theoretically relevant to the research question. These dimensions included age, profession, sector, national and cultural background, and prior engagement with technology. The aim was to surface the widest possible range of ideological articulations within a feasible number of interviews (Patton, 2002). Snowball sampling was used to extend the pool through respondent referrals, particularly where these referrals broadened rather than narrowed the diversity of the

sample. The outreach was conducted through the researcher's personal and professional networks via WhatsApp, Instagram, and LinkedIn. This was supplemented by direct invitations to individuals identified for their relevance to the research question.

The final sample comprises twenty-four interviews involving twenty-five individual respondents. Twenty interviews were conducted with consumer respondents drawn from professional, entrepreneurial, academic, and student backgrounds. One of these was a joint interview with two participants from Hong Kong, who are reported as separate respondents in the analysis. Four interviews were conducted with expert respondents. Their direct engagement with humanoid robotics, artificial intelligence, or technology-driven business transformation positioned them to contextualise the consumer narratives and to comment on developments at the frontier of the phenomenon. The expert interviews were not used as a separate triangulating dataset. They were integrated into the analytical corpus, on the understanding that experts also speak as ideological subjects whose narratives can be read for their articulation of technology meanings.

Respondents range in age from their early twenties to their mid-fifties. Their professions span primary school teaching, hospitality management, law, capital markets, computer forensics, molecular biology, nutrition, music, marketing, and diplomacy, alongside senior executive roles in insurance and industrial construction. Their geographic backgrounds cover Austria, Germany, Switzerland, Italy, Spain, France, Luxembourg, the United Kingdom, Bulgaria, Poland, Romania, Sweden, the United States, and Hong Kong. Three respondents are parents of young children. Several work in roles directly exposed to potential automation, and several others occupy creative or interpersonal professions for which automation is a more relevant prospect. The four expert respondents are: the founder of Tesla Owners of Silicon Valley; the Chief Information Officer of an Austrian industrial construction group integrating robotic systems into its operations; a marketing and artificial intelligence professor working with generative AI since 2019; and a doctoral researcher at the University of Vienna who lives and works alongside a humanoid research prototype. A summary table is provided as Table 1.

Table 1 Respondent composition

Name	Age	Profession	Location	Gender	Group
Alexander Ietan	—	CIO, industrial construction group	Austria	M	Expert
Emily Kate Genatowski	34	Doctoral researcher; lives with humanoid prototype	Austria/USA	F	Expert
John Stringer	40	Founder, Tesla Owners of Silicon Valley	USA	M	Expert
Martin Reisenbichler	—	Marketing professor and AI specialist (gen AI since 2019)	Austria	M	Expert
Pseudonym	Age	Profession	Location	Gender	Group
Adrien	Late 20s	Geopolitical surveillance focus	Luxembourg	M	Consumer
Anya	Early 30s	Marketing manager	United Kingdom	F	Consumer
Bruno	Early 20s	Bachelor student (data and AI)	Austria	M	Consumer
Charlotte	Late 20s	Corporate lawyer	Austria	F	Consumer
Elena	Early 30s	Finance, controlling, and reporting	Kosovo	F	Consumer
Felix	Early 20s	Master's student, marketing	Austria	M	Consumer
Henrik	Early 30s	Professional (work-displacement concerns)	Armenia	M	Consumer

Name	Age	Profession	Location	Gender	Group
Ho	20s	Computer forensics, financial-crime investigation	Hong Kong	M	Consumer
Isabel	Late 20's	Molecular biologist (catalyser industry); mother	Sweden	F	Consumer
Karin	Mid-20s	Public health and psychology researcher	Germany	F	Consumer
Liliana	Late 20s	Big Four professional	Italy	F	Consumer
Lorenz	40s	Self-employed nutritionist and personal trainer	Austria	M	Consumer
Lukas	Late 20s	Computer scientist, fintech programmer (MSc AI)	France	M	Consumer
Markus	45	CEO, Swiss office of international insurance broker	Switzerland	M	Consumer
Matteo	Early 30s	Marketing professional (AI-generated media)	Italy	M	Consumer
Maya	20s	Capital markets (prev. Bloomberg)	Hong Kong	F	Consumer
Nadia	Mid-20s	Marketing campaign manager; one child + expecting	Bulgaria/Poland	F	Consumer
Petar	Mid-20s	Music student and professional musician	Greece	M	Consumer
Sofia	Late 20s	Primary school teacher (international school)	Ireland	F	Consumer

Name	Age	Profession	Location	Gender	Group
Stefan	Late 20s	EU delegation diplomat	Germany	M	Consumer
Tomáš	Mid-20s	Master's student, marketing; entrepreneur in 3D printing	Austria	M	Consumer

Full per-respondent profiles of the experts are presented in Appendix E.

The target sample size was set in advance at approximately twenty consumer interviews plus a small number of expert interviews. This was in line with established qualitative practice for in-depth interpretive studies of comparable scope (Guest, Bunce, & Johnson, 2006; Kozinets, 2008, who built his foundational model from six respondents). The decision to conclude data collection at twenty-four interviews was guided by two considerations. First, the intertextual analysis had begun to yield clearly recurring patterns across heterogeneous respondents, a working indicator of conceptual saturation in interpretive research (Saunders et al., 2018; Bowen, 2008). Second, the inclusion of further respondents drawn from the same purposive frame was unlikely to surface ideological articulations categorically different from those already documented. The researcher acknowledges that conceptual saturation in interpretive research is never absolute. Further sampling, particularly across cultural contexts not represented here, would certainly yield additional refinements.

3.4 Data Collection

The interviews were conducted between February and March 2026. All but one took place via Google Meet; the remaining interview was in person. Interview duration ranged from approximately 30 minutes to 90 minutes, with most falling between 45 and 75 minutes. The variation reflects the natural rhythm of each conversation and respondents' willingness or capacity to elaborate, rather than any restriction imposed by the researcher. All interviews were conducted primarily in English, with brief excursions into German where this aided respondent expression. Quoted excerpts in non-English passages have been translated for analytical use, with the original retained in the research files.

The interviews followed a semi-structured guide organised into five thematic sections (the full guide is reproduced in Appendix A). The guide opens with a warm-up that

establishes recent encounters with humanoid robots in the respondent's media environment, then moves through purposes and applications (with deliberate attention to the home as a context), trust, social and labour implications, and a closing section eliciting overall sentiment and information sources. Where respondents produced thin or terse answers, follow-up probes invited elaboration, example, or counterexample. Where respondents produced rich or unexpected material, the guide was set aside in favour of pursuing the emerging line of thought, in line with the principles of in-depth interpretive interviewing (McCracken, 1988; Kvale & Brinkmann, 2009).

All interviews were audio-recorded with the respondent's explicit consent. Recordings were subsequently transcribed verbatim and lightly cleaned for readability. The Transcripts were stored in the researcher's working corpus under the respondent's first name. In the case of expert respondents who agreed to be identified, their full name was used. Identifying information for non-expert respondents has been removed or generalised in the final thesis text.

3.5 Ethical Considerations

The study was conducted in accordance with the principles of voluntary participation, informed consent, anonymity for non-expert respondents, and secure data handling. At the start of each interview, the researcher introduced the project as a master's thesis concerned with public perceptions of humanoid robots and artificial intelligence. The researcher indicated that the conversation would be recorded for transcription purposes and invited the respondent to provide explicit verbal consent on the recording. Respondents were informed that their participation was voluntary and that they could decline to answer any question. Non-expert respondents were further informed that they would not be identified by name in the published thesis. The deliberate choice not to brief respondents on the specific theoretical framework guiding the analysis was made to preserve the spontaneity of their articulations. Respondents were nonetheless informed of the broader topic and purpose of the conversation, which is consistent with the standards of informed consent in interpretive consumer research (Belk, 2017).

The four expert respondents were interviewed in their professional capacities and consented to be identified by name. Audio recordings and full transcripts are retained in the researcher's secure working files and will not be circulated beyond the supervisory chain. The closing of each interview offered the option of receiving the completed thesis on request. The study did not involve vulnerable populations,

sensitive personal data, or interventions that posed any foreseeable risk of harm to respondents.

3.6 Researcher Positionality and Reflexivity

In interpretive research, the researcher is the primary instrument of data construction and analysis. The methodological soundness of the study therefore depends in part on the researcher's reflexive engagement with his own positionality. The author approaches humanoid robotics from sustained personal and professional engagement with the field. He is a member of Tesla Club Austria, active in Lumos, an Austrian student technology organisation, and has spent considerable time in technological enthusiast communities. Through these affiliations he has developed familiarity with the discourses, expectations, and knowledge that circulate within humanoid robotics communities.

This positionality carries both advantages and risks. On the one hand, sustained engagement with the field has produced an interpretive sensitivity to respondent narratives that a researcher coming to humanoid robotics cold would have to work harder to develop. On the other hand, it carries the risk of unwitting alignment with technological enthusiast framings of the phenomenon. Hence, three concrete measures were taken to manage this risk. Firstly, the interview guide was deliberately constructed to invite expression of resistance and scepticism, with multiple probes oriented to surfacing concerns. Secondly, the analytical procedure described in section 3.7 deliberately held two approaches in tension: deductive engagement with existing theory and inductive openness to new patterns. When an interpretation seemed only to confirm the researcher's prior assumptions, the procedure required a return to the data. Thirdly, an explicit search was conducted, during the cross-case synthesis cycle, for respondent material that contradicted the researcher's initial intuitions. Several of the patterns identified in Chapter 4, including the strength of the Human-Reservation orientation and the velocity-anxiety dimension, emerged through this counter-intuitive search.

3.7 Data Analysis

The analytical procedure followed a hermeneutic part-to-whole interpretation in the tradition of Thompson (1997) and Spiggle (1994). It was adapted for this study that aims to both test and possibly extend an existing model. Therefore, it proceeded in five overlapping cycles.

The first cycle was an intratextual reading. Each transcript was read in full as a complete narrative, with the aim of reconstructing the respondent's position as a coherent (though often internally contradictory) ideological articulation. The researcher composed brief interpretive notes for each respondent, recorded the dominant ideological registers, the principal contradictions, the emotional tonality, and any narrative episodes through which the respondent worked through the topic. The same words read differently in isolation than they do in a respondent who, two minutes later, recoils from the same prospect.

The second cycle was a deductive coding pass against Kozinets' (2008) model. Across the full corpus, the researcher coded passages for evidence of the four ideological elements (Techtopian, Green Luddite, Work Machine, Techspressive) and the inner contradictions Kozinets identifies within each. Passages were coded as articulations of an element when their nodal structure clearly invoked one of Kozinets' supreme goods (Progress, Nature, Economic Growth, Pleasure). Special attention was paid to within-respondent movement between elements, which Kozinets identifies as the model's most distinctive empirical signature, and to the triggers that prompted such shifts.

The third cycle was an inductive open coding pass. Working again across the full corpus, the researcher coded for themes that did not map cleanly onto Kozinets' four elements, treating these excess meanings as theoretically productive rather than as noise. This was primarily done to identify the model's blind spots and possible extensions.

The fourth cycle was an intertextual synthesis. The patterns from the deductive and inductive cycles were tested across the full corpus to identify which articulations were widely shared, which were concentrated in particular subgroups, and which appeared in surprising combinations. Attention was paid to configurations (the recurring couplings of ideological elements that gave individual accounts their shape) and to cleavages (the points at which the corpus divided most sharply). This work produced twelve cross-cutting patterns, six of which structure Chapter 4 while the remaining six are documented in Appendix F.

The fifth cycle was a theoretical reflection on the relationship between the findings and Kozinets' (2008) model. Each pattern was examined for whether it affirms, challenges, or extends the model in the present technological moment. These reflections are developed in Chapter 5.

The analysis was conducted by the researcher through close, repeated reading rather than through dedicated qualitative data analysis software, on the judgement that a corpus of twenty-four transcripts is small enough to hold in interpretive view as a whole. A large language model (Claude Opus 4.7) was used as a supporting tool during analysis, for tasks such as surfacing candidate themes and cross-checking coding consistency. All interpretive judgements, code assignments, and theoretical claims were made and verified by the researcher. The researcher maintained an analytic journal throughout, recording emerging interpretations, doubts, and revisions (Spiggle, 1994).

3.8 Methodological Limitations

The sample, while heterogeneous along multiple theoretically relevant dimensions, is not representative of any specific population. Due to the recruitment from within the researcher's extended network, respondents are disproportionately drawn from European urban professional and student backgrounds. Although several respondents bring non-European cultural backgrounds to the interview, the sample underrepresents respondents from outside the European, North American, and East Asian cultural spheres. Additionally, all coding and interpretive work was conducted by a single researcher, without inter-coder triangulation. This is mitigated by the procedural transparency described above and by the explicit reflexive practice, but it still remains a constraint. Most importantly, the study captures a particular and rapidly evolving historical moment in the public emergence of humanoid robotics. Respondent narratives are inevitably shaped by the specific media coverage and public discourse current in early 2026. Had the researcher conducted this thesis twelve months earlier or later, respondents would almost certainly yield different ideological articulations. Lastly, the absence of pre-stimulus material is both a strength and a limitation. As a strength, it preserves the spontaneity of respondents' existing mental models. As a limitation, it constrains the analysis to the imaginative resources respondents already possess and cannot speak to how those resources would shift if respondents were exposed to direct experience with the technology. Every limitation is taken up again in Chapter 6, where their implications for the generalisability and durability of the study's contributions are considered.

4 Findings

This chapter presents the empirical findings of the study. It is organised around six cross-cutting patterns that emerged from the analysis of the twenty-four interview transcripts described in Chapter 3. Each pattern is a recurring articulation, contradiction, or structural feature of the consumer narratives in the dataset. Together, the six patterns form the empirical material against which Kozinets' (2008) ideological field model will be tested in Chapter 5. Six further patterns that emerged in the analysis but bear less directly on the model extensions are documented in Appendix F. The patterns are not presented in order of explanatory weight, but sequenced to allow earlier interpretive moves to inform later ones. This chapter opens with the most general feature of the corpus, the qualified optimism that almost every respondent ultimately settles into, and closes with a synthesis of the configurations and cleavages that the patterns generate.

To help orient the reader through the patterns, Figure 2 previews the revised model that Chapter 5 develops. It is introduced here in its three constituent layers. Panel (a) presents the field plane: the four original elements Kozinets (2008) specifies, together with the emergent fifth element this corpus surfaces. The fifth element orbits the field as a satellite element, rather than occupying a corner of it. Panel (b) presents the two substrates, which act as structural conditions that envelop the field as a whole. Panel (c) presents the velocity dimension, along which the field itself is in motion. Each of the following patterns feeds one of these layers, and the relation is indicated at the close of each section. The assembled model, together with the full argument for each of its components, is presented in section 5.4.




Figure 2 The revised ideological field model introduced in three layers.
Elements and semantic relations adapted from Kozinets (2008)

4.1 The Qualified Optimism Paradox

The single most pervasive feature of the corpus is also its most analytically perplexing. Almost without exception, respondents conclude their interviews by declaring themselves optimistic about the future. They do so even when the substantive content of the preceding conversation has dwelt, sometimes extensively, on labour displacement, surveillance, the erosion of human connection, the concentration of corporate and political power, the prospect of cognitive decline, and outright dystopian scenarios up to and including the so-called Terminator outcome. Of the twenty-five respondents who answered the closing optimism question, twenty positioned themselves explicitly as optimistic, three as neutral, one as “half and half,” and one as honestly unable to say. Not a single respondent declared themselves pessimistic about the future, despite many having articulated pessimistic content throughout. This finding is not about respondents’ underlying sentiment. Instead, it illuminates about the narrative architecture of contemporary technology talk.

Respondents typically open with mixed or curious affect, develop a series of concerns through the middle of the interview, and arrive at the closing question with a routine formulation of optimism that reads as a moral resolution rather than an empirical conclusion. Sofia, a primary school teacher, summarizes the logic plainly: “If you’re pessimistic about it, that doesn’t do anyone any good. So yes — stay positive.” Elena arrives at a structurally identical position: “If you’re pessimistic about the future, your present is miserable. How can you live your daily life thinking the future is going to be terrible?” Even respondents who have spent the bulk of the interview expressing significant concerns end in the same register. Henrik, whose negative articulations are among the strongest in the corpus, concludes: “We will see what happens. We have to be happy somehow, we have to go on, do the best of what we do best in our lives. That’s all we can really do.” Alexander Ietan, the CIO of an industrial construction group integrating robotic systems, produces perhaps the most analytically revealing version. He declares himself optimistic, but immediately qualifies the optimism with an escape fantasy: “Maybe I’ll end up living somewhere in Alaska without a robot. That might also be nice.” The fantasy contradicts the very technological future he has spent an hour endorsing.

This pattern is not denial or social desirability bias. It should be read as the rhetorical operation of an ideological field that does not currently provide respondents with a viable subject position from which to declare themselves pessimistic about technology

in general. Optimism functions, in the closing register of these conversations, less as a claim about the world than as a signal of one's capacity to inhabit it. It is a performance of the kind of psychological and moral disposition the contemporary moment requires. The qualified optimism is qualified precisely because the substantive concerns have not been resolved. They are put aside, so that the speaker complies in accordance with current societal narratives.

In relation to Kozinets' (2008) model, this pattern affirms the Techtopian element's continued cultural dominance. The supreme good of Progress remains the default to which respondents return when forced to summarise. At the same time, the pattern reveals something the model does not foreground. Often, the dominance appears to be performative rather than substantive. Respondents inhabit Techtopian closure even when the bulk of their narrative articulates Green Luddite, Work Machine-pessimistic, or fundamentally Mick-and-Fournier-paradoxical content. The Techtopian element, in this corpus, functions less as a stable ideological position than as a terminus of permitted speech. This is a finding to which Chapter 5 returns.

The pattern this section documents can be located against the available population-level evidence on adjacent attitudes. The 2024 Special Eurobarometer 554 found that 55% of EU citizens expect AI to positively shape daily life over the next 20 years, against 35% who anticipate negative effects, and that more than 60% view robots and AI at work favourably (European Commission, 2025; Verian Group, 2025). Pew Research Center's parallel surveys of US adults register substantially more ambivalence: 52% of Americans report feeling more concerned than excited about AI's growing role in daily life (a 14-point rise since 2022), and only 17% of the general public believes AI will have a very or somewhat positive impact on the United States over the next 20 years (Pew Research Center, 2025).

The European population appears closer to qualified optimism than the American, and the gap between expert and lay opinion is particularly stark in the US. The closing-optimism architecture of the present interviews therefore overstates positive sentiment relative to representative survey means. Three features of the sample plausibly amplify the optimistic tilt: the European concentration of the sample, the educational and professional composition (Pew finds higher-income, more-educated respondents foresee positive impacts more often), and the dialogic interview context itself, in which the rhetorical pressure toward a coherent closing posture is structurally greater than in an anonymous survey instrument.

4.1.1 Coping Strategies and the Architecture of Optimism

The qualified optimism pattern admits of a further interpretation that locates it within an established framework in consumer research. Mick and Fournier (1998), in the most theoretically generative empirical study of consumer-technology paradox in the discipline to date, document eight paradoxes that technology consumption activates. They identify a parallel set of coping strategies through which consumers manage the affective load these paradoxes generate. They identify the cognitive-emotional work of paradox management as central to the lived experience of technology in consumer life.

The present corpus, read through this framework, suggests that the qualified optimism documented above is not the absence of distress but a specific form of its management. It is a coping strategy that operates at the level of rhetorical closure. The reason is that the technology in question is not yet present in the lives of the great majority of respondents, and the more behavioural strategies Mick and Fournier identify (mastery, partnering, neglect) are not yet available. What is available is the interpretive resource of a closing register, and that resource is mobilised consistently across the corpus.

Several distinct coping strategies are nonetheless legible in the substantive content of the interviews, and they map onto the Mick and Fournier (1998) typology. Pre-emptive distancing, the strategy of avoiding the technology before encounter, is articulated most fully by Liliana, who says: "I would just rather see how others use it for a while, gather some advice, and then after a while embrace the change." Isabel articulates the same strategy in the form of declared refusal regarding specific use cases. Confrontation through disciplined use, that is, engagement with the technology under self-imposed constraints designed to mitigate its perceived risks, is articulated most fully by Tomáš. He frames his deliberate practice ("I've also made it a point to use AI as deliberately and sparingly as possible — to provoke thoughts, not substitute them. I use it to check an article I've written, not to write one for me") as a personal protocol against the cognitive atrophy he anticipates. The position is supported by Elena's parallel discipline of cross-verifying AI outputs against independent research and by Sofia's emphasis on guarding the relational practices that the technology would otherwise displace. Mastery, the strategy of accumulating sufficient technical knowledge to neutralise the anxiety the technology generates, is articulated most fully by Bruno. His three-layer epistemic strategy (university training, then research papers, then online materials) reads as a deliberate construction of interpretive competence that, in his own framing, allows him to "categorise" incoming information that would

otherwise overwhelm. Partnering, the most behaviourally substantive of Mick and Fournier's strategies, is articulated only by Emily Kate Genatowski, whose direct cohabitation with a humanoid prototype constitutes the operational form of the strategy.

The principal extension this corpus suggests to Mick and Fournier's framework concerns rhetorical closure as a coping strategy in its own right. The 1998 framework treats coping as a set of cognitive and behavioural responses to the lived encounter with technology. Under conditions of anticipated rather than experienced technology, a further coping mechanism appears to become available: the construction of a closing optimistic stance whose function is to discharge the affective load of substantive concern that has accumulated through the interview. It is not mastery, since the speaker has not acquired interpretive competence; not avoidance, since the speaker has engaged at length with the concerns; not partnering, since no technology is present to partner with. It is a discursive coping strategy directed at the speaker's own affective state, performed under the social conditions of the interview.

The closing optimism appears to function not mainly as a commitment to the supreme good of Progress, but as a discursive instrument through which respondents discharge the unresolved affective load of paradox engagement. The substantive Tectopianism that section 5.1 affirms (Alexander Ietan, John Stringer, Emily) coexists with this performative variant, and the analytical task is to distinguish between them rather than conflate them.

4.2 Surveillance as Master Fear

Surveillance is the most prevalent concern among respondents in the corpus. Where many of the other patterns divide respondents or amplify their existing positions, the fear of surveillance unites them. Respondents from across the ideological spectrum, from the most enthusiastic Tectopians to the most resistant Green Luddites, converge on data privacy and the unaccountable transmission of household information as the single most pressing concern raised by humanoid robotics. Respondents who agree on almost nothing else, agree on this.

Alexander Ietan, who is otherwise among the most enthusiastic Tectopian-Work Machine voices in the corpus, identifies surveillance as the issue he most wants future regulation to address: "The much more urgent dimension is data privacy. Humanoid robots will see a great deal of your private life. They will collect data and send it somewhere. We absolutely have to regulate that." Also Charlotte, articulates the same concern: "I'm probably more cautious than most about privacy. I don't want a robot

filming and hearing everything in my home.” Liliana identifies the visual dimension specifically: “Robots running on cameras — that’s an enormous amount of sensitive visual data. Where does it go? Who stores it? Is it a government database?”

The convergence holds even where respondents disagree about virtually everything else. Matteo, endorses humanoid adoption and even frames the robot as a “slave” in functional terms, however, even he has his concerns about the trustworthiness of the manufacturer on data-handling grounds: Tesla yes, Chinese suppliers initially no. Even respondents who prefer to live with a robot (Emily Kate Genatowski, who has done so for over a year) anchor much of their professional work on regulatory gaps in data privacy and security, identifying these as the elements of the field that most urgently require legal development.

Several respondents connect the surveillance concern to a temporal observation that is not commonly registered in earlier consumer technology research. The home, they note, has historically been the boundary of legitimate observation. The smartphone in the pocket, the smart speaker on the shelf, and the camera doorbell in the hallway have already eroded that boundary. The humanoid in the corner of the living room would, in their reading, complete its dissolution. Lukas captures the underlying suspicion succinctly: “The closest thing most people have to a humanoid robot at home right now is Alexa or Siri... and my problem with those is they’re constantly listening and nobody fully knows what happens with that data. A humanoid robot would trigger similar privacy concerns.” The humanoid is not, in this reading, the source of the concern. It is the end of a longer trajectory that respondents have been observing and that they understand, perhaps for the first time, to be approaching soon.

Kozinets’ (2008) identified surveillance as an inner contradiction of the Techtopian element. Yet, he did not develop this into a distinct structural feature of the field. In contemporary times surveillance functions as a substantive nodal point of resistance that any technology adoption narrative now has to negotiate, regardless of its ideological starting position. The contribution of the present study, on this point, is to surface a fear that sits underneath the four elements rather than within any one of them. This is a finding to which Chapter 5 returns in detail.

4.3 Geopolitical and Corporate Origin as Trust Proxy

This pattern concerns the question of who, exactly, is making the robot. For an unexpectedly large share of respondents, the question of trust in humanoid robotics is not primarily about the technology itself but about its provenance, both national (which country produced it) and corporate (which firm and under what value system). This shifts the analytical centre of the trust question away from the device and toward the geopolitical and ideological context of its production and roots.

The geopolitical dimension is articulated greatly by respondents who work in or adjacent to public-sector environments and by respondents whose own national identities sit at the boundary of the dominant blocs. Stefan, whose professional life is structured around precisely such distinctions, names the established pattern from his work: “You always have the same actors — China, Iran, Russia on one side, and then the US, which has increasingly become unreliable in ways no one would have predicted five years ago.” Alexander Ietan expresses his European perspective: “As Europeans, we need European solutions. A lot of companies will not hand all their production data to a US or Chinese company.” He invokes the US Cloud Act explicitly as a structural reason for European data sovereignty, and identifies the German Telekom sovereign cloud infrastructure as the right direction of travel. Adrien locates the specific concern in China’s social credit infrastructure: “Countries like China, where you already have a social credit system and a huge surveillance infrastructure — that’s a specific concern. That technology could be used to limit and control people in ways that are genuinely frightening.”

The corporate dimension overlays the geopolitical one and, for several respondents, partially supplants it. John Stringer, whose location at the centre of one of the most visible technology communities makes his framing analytically significant, places the question of the manufacturer’s ideology at the centre of the trust calculation: “It matters enormously who is creating the AI, who is creating the robots, what ideology is informing them, and how that reflects in the decisions these robots make.” He extends the framing through a brand analogy: “There are companies I would never buy a robot from because of how they think and the information they’ve put out — we saw it during COVID, with misinformation being spread. That’s the ideology question that matters.” Bruno, whose direct technical engagement with the AI laboratory landscape is the most current of any respondent, draws a precise corporate distinction within a single national context: “The two big AI labs right now are OpenAI and Anthropic. I would trust Anthropic significantly more than OpenAI. Maybe it’s because of things like Sam Altman going on camera saying they will never introduce ads... and

then a year later introducing ads. If a company has a track record of sticking to what they say and not switching positions or being involved in scandals, that builds trust and credibility.” Tomáš develops the structural version of the same argument. Any AI model carries the biases of its training data and its developers, and those biases will be transferred to users. The question of provenance is therefore not extrinsic to the technology but constitutive of it.

Several respondents combine the geopolitical and corporate registers explicitly. Stefan, after laying out the country-level analysis, adds a layer that he finds perhaps more important: “It’s not just countries — the multinational corporations and certain super-rich individuals also operate with different objectives than states. Countries are at least somewhat legible in terms of their motivations. With private entities, the number one aim is presumably profit, but for people with hundreds of billions, making more money isn’t really the objective. So then the question of what they’re actually pursuing becomes more interesting and less answerable.” Bruno draws the same connection from a different direction, identifying the concentration of power at the intersection of major AI labs and political administrations as his single greatest concern about the field’s trajectory.

A small number of respondents (Martin Reisenbichler, Charlotte) qualify the importance of geopolitical origin by foregrounding security risks (hacking, botnet incorporation) over national provenance. Even they accept the premise that the institutional context of production strongly affects the trustworthiness of the device. Therefore the minority position in this corpus is not that origin doesn’t matter... it is that origin matters less than security architecture.

This pattern substantially extends the original model. The 2008 framework was developed in a moment when the dominant ideological articulations of personal computing did not significantly thematise the geopolitical or corporate-ideological provenance of the technology. The technology was largely American and the corporate identities involved were not yet inscribed with the political influence they have since acquired. Hence, consumer narratives could afford to treat the device in relative isolation from its production context. In contrast, the present corpus demonstrates that consumer narratives about humanoid robotics cannot do so. Provenance is no longer a background fact but a foreground variable in the trust calculation, and any contemporary mapping of the ideological field must register it. This is one of the clearest extensions the data demands of the model, and Chapter 5 develops the theoretical implications.

4.4 Velocity Anxiety: Pace as Ideological Provocation

This pattern shifts the focus from what the technology is to how fast it is arriving. A consistent strand across the corpus articulates an anxiety that is not principally about the consequences of humanoid robotics. The anxiety, is about the rate at which the technology, and the broader artificial-intelligence complex of which it forms part, is developing. The pace is itself the provocation.

Henrik gives the pattern its most introspective formulation. When asked if keeping up with AI discourse is stressful, he responds with an honesty that the rest of the interview structures into an extended self-reflection: "There's a term... FOMO, fear of missing out. I only learned about that term one or two months ago. And somehow, even though I resist it, I find myself on X every day seeing new updates, new models, new capabilities. I ask myself: why? What is my benefit from this? It's fear of missing out, and I'm actively trying to let that go." He then offers an observation about ideological in-group sorting that is striking in its sociological precision: "If you're inside this AI and technology bubble, and you compare yourself to people who are not in it, they seem happier. Not because they're avoiding hard things, but because they're not constantly confronted with questions they can't answer." The pace of development, in this reading, is not solely technological but epistemic. It generates more questions than the average human can deal with, resulting in a state of chronic incompleteness.

Tomáš, reinforces this sentiment: "I don't think we're capable of coping with the speed of technological advancement. It's not just my opinion; all the data suggests that we're not." He grounds the claim in evolutionary psychology and in the documented mental health correlates of social media adoption, framing the pace as a structural mismatch between the rate of technological change and the rate at which human beings can adapt cognitively and emotionally. Bruno, locates the same difficulty at the professional level: "Keeping up with all the newest developments is basically a full-time job. There are so many companies building things... so much development that it's nearly impossible to track. Even with my thesis — I try to get my resources together and write, and then two weeks later there's a new research paper in the same field and I have to factor that in. I don't know where to stop." Liliana, a deliberate slow adopter, articulates the inverse posture as a coping strategy. She would let others test the technology first, "gather some advice, and then after a while embrace the change,".

Interestingly, the pace is not univocally negative across the corpus. A second strand articulates the same temporal acceleration as a source of exhilaration rather than

anxiety. Markus, a Swiss insurance executive, frames it as rare biographical privilege: "I'm genuinely happy to be alive right now because we're witnessing an incredible change." John Stringer elaborates further: "There has never been a cooler or more exciting time to be alive. In our lifetime, we've experienced three industrial-ish revolutions — the internet boom, the smartphone boom, and now the AI and AI robotics boom. We've lived through the greatest technological changes in the shortest time period of any generation in history." Emily Kate Genatowski, offers the fullest articulation of pace-as-privilege: "We are standing at the precipice of a completely new world. We are incredibly lucky to have been born at this time, to have these interests, to have backgrounds that have led us to this field right now."

The same temporal phenomenon, is being described by different respondents as the source of two opposed affective experiences. The anxious narration treats the pace as overwhelming, exhausting, or even incompatible with healthy human functioning. The excited narration treats the same pace as the historical condition of meaningful agency and a rare opportunity to participate in a transformative moment. The difference is not between respondents who agree about the pace but disagree about its valence; both sides agree about the empirical fact. It is about whether the human subject is positioned to ride the velocity or to be carried by it. That positioning is a structural feature of the respondent's broader ideological orientation, not a separate variable.

Kozinets analyses ideological positions on technology in relatively static terms. Respondents inhabit and move between elements, but the model does not theorise the temporal modality of the technological field as itself an ideological object. The present corpus suggests that in a moment of perceived acceleration, the pace itself becomes one of the things consumers must take a position on, and that the position they take is partially constitutive of their broader ideological articulations. Chapter 5 develops the implications.

4.5 Human-Made Value as Resistance Ideology

This pattern names an ideological articulation that several respondents developed: the proposition that human-made objects and services might become more valuable, not less, as artificial production becomes more capable. Imperfections and touch of human labour are identified as a residual premium that the technologically saturated future will preserve rather than dissolve. The articulation resists ideologically the wholesale automation narrative. It is not nostalgic refusal, but anticipation of a market segment and a cultural domain that will remain reserved for the human.

Tomáš offers his perspective: "If a humanoid robot makes a bag and a person makes the same bag, the robot's will be flawless. The human's will have two or three imperfections. People will buy the human-made one and pay more for it, because of human connection. I think that's true for photography, apparel, any artistic form. AI can make nicer photos, but they're AI-made photos. We long for human connection, even in the objects we own." His argument is equally economic, aesthetic, and ethical. The human-made bag commands a price premium because the imperfection it carries is read as evidence of the human producer's presence. The AI-generated photograph might be technically superior but ontologically depleted. Consumers precisely long for the human touch behind a photo, a value that AI cannot supply.

Isabel develops a structurally identical argument from a different starting point, her own practice of ballet, and extends it into a warning about cultural standardisation: "With AI, everything starts to look the same. Everyone writes in the same way, everything feels the same. People are losing their singularity." The remark is interpretively double-loaded. Singularity, in Isabel's usage, refers to individual human distinctiveness, the irreducible particularity that AI-mediated production tends to homogenise. The same word, in the technical literature on artificial intelligence, refers to a hypothetical point of self-improving machine intelligence beyond human comprehension. The collision of meanings is itself a finding. The technological discourse and the ideological resistance to it have, in this respondent's articulation, converged on the same word from opposite directions.

Nadia develops the most historically expansive version of the human-made-value argument through an exposition of art history: "Throughout history, creative fields have always saturated and then transformed." She traces the movement from classicism to romanticism to free jazz to atonal abstraction. She frames each transition as a response to the saturation of the preceding form, and predicts a comparable response to AI-generated content: "Websites written entirely by robots will saturate people. Something new will emerge that humans do differently." She closes the exposition by anticipating the return of philosophy as a viable profession: "sitting in cafés, writing for newspapers, thinking. We lost that in the last fifty or seventy years. I think we can get back to that kind of work." This historical narrative positions human-made cultural production not as a residual preserve but as the next movement, generated by the saturation that AI itself will produce.

Several respondents articulate similar versions of the same pattern. Sofia returns spontaneously to the value of human encounter in passing: "Every person sitting next to you has an entire world inside them, a whole life and experience, and you have an

opportunity right now to connect with them. I just hope that as technology advances, people are still able to have true human connections outside of it.” Stefan develops the same intuition through the specific case of the cleaning lady at his workplace, whose job, he acknowledges, could be done by a robot, but whose human presence is part of what makes the workplace enjoyable and whose loss would constitute a real subtraction from his daily life. Even Emily Kate Genatowski anchors her optimism in the irreducibility of the human. She tells the story of a friend with severe dyslexia who, in a game designed to require the player to convince an AI that they were human, was flagged as human on the very first message because of his spelling. “There’s something so beautiful about that,” she says, “these very humanistic elements, these things that resist categorisation.”

The pattern is a positive articulation of what the domestic-sphere pattern (Appendix F.1) negatively defends. Where the home is the spatial locus at which the human is held in reserve, human-made value is the productive and aesthetic locus at which the same reservation is articulated. The two patterns are mutually reinforcing. The home is sacred because it is where humans live, and human-made things command a premium because they carry the essence of that same humanness the home is meant to protect.

In Kozinets’ (2008) terms, this pattern is not properly captured by any of the four original ideological elements. It is not Green Luddite. It is not articulated as opposition to technology, but as anticipation of a market and cultural arrangement in which the human and the technological coexist with differentiated value. It is not Techspressive. It does not articulate technology as a vehicle for individual self-expression, but the human artefact as the bearer of authentic value. It is therefore an emergent ideological articulation that the model needs to accommodate. In this position the supreme good is neither Progress nor Nature nor Productivity nor Pleasure, but the irreducibility of the human itself. Chapter 5 develops the theoretical implications.

4.6 Cognitive Atrophy as Emergent Anxiety

Across several respondents, anxiety is articulated about the potentially devastating cognitive consequences for users of widespread artificial-intelligence and robotic adoption. This fear is based in the slow erosion of human capability under conditions of progressive cognitive offloading. The user, in this articulation, is not threatened by the technology, but rather diminished by their own use of it.

Elena shared an observation: “I read something recently suggesting that up until millennials, average intelligence was growing steadily, and now the curve is going

downward for the first time. I think that's partly because young people are using AI to finish their homework rather than using it to genuinely learn." She extends the observation into a principle that she applies to her own use of the tools: "I use ChatGPT and other AIs daily, but I use them to get to the core of information and then verify it... I never rely on it exclusively." The deliberateness of the practice, the active resistance to over-reliance, is itself a signature of the pattern. Respondents who articulate cognitive atrophy concerns typically also describe the personal practices through which they attempt to mitigate it.

When Tomáš was asked if he would want a humanoid robot at home, he responds with an argument that connects technology use, cognitive function, and personal practice into a single thesis: "I've also made it a point to use AI as deliberately and sparingly as possible — to provoke thoughts, not substitute them. I use it to check an article I've written, not to write one for me. That comes from a belief in evolutionary psychology — I don't think we can evolve quickly enough to collaborate healthily with these technologies." He goes on to cite empirical work: "There's actually research — I recall one study that scanned the brains of people who wrote essays themselves versus those who used ChatGPT, and over time their ability to write those essays declined the more they relied on AI. If a language model can already do that, imagine what a physically present AI would do to that effect." The argument is understandable. AI use offloads cognitive work. Offloaded cognitive work atrophies. Atrophied capability cannot be recovered on demand. The embodied form of AI (humanoid robots) intensifies the offloading and therefore the embodied form intensifies the atrophy.

Isabel recognizes the pattern in her student context: "I've heard about a lot of students using AI to write their theses, and I find that insane. You chose university because you want to learn, and if you use something to do it for you, you never learn... If you write your thesis using AI, then later in your career you won't be able to write a paper without it." Henrik returns to the same argument from a physiological metaphor: "The brain, like the muscles in the body, needs training. If you give your brain new things to learn, it stays sharp. The same with muscles — train them, keep them. And I think about this when it comes to AI and robots... Using these tools for things you already know how to do is not a net gain, it's a net loss." He extends the argument into a behavioural critique that connects directly to the home-robot debate analysed in Appendix F.1. A robot that does the laundry frees up time, but the time freed up is, in his observation, statistically far more likely to be redirected toward additional screen consumption than toward any of the higher-order activities that automation's defenders typically invoke.

Lorenz adds the historical-philosophical register, citing Theodore Kaczynski's arguments about technological dependence and explicitly disclaiming Kaczynski's methods while endorsing the underlying analytical concern. He reports having recently bought a paper notebook for project planning, on the considered view that the digital alternatives, were eroding capacities he wished to preserve.

The cognitive atrophy concern is not articulated by the most resistant or anti-technological respondents. It is articulated, by respondents who are themselves heavy users of the relevant tools and who frame their concern in self-implicating rather than externally accusatory terms. The pattern is therefore not a Green Luddite formation in the original Kozinets (2008) sense. It does not articulate technology as the enemy. Instead, it is a reflexive position: the user is the locus of the problem, and the technology is creating conditions under which the user's own self-discipline is the only available defence.

This pattern significantly extends the original model. The 2008 framework identifies the inner contradictions of the Techspressive element as the fragility of technologically mediated pleasure, yet, does not theorise cognitive degradation as an inner contradiction of any of the four elements. The present corpus shows that cognitive atrophy is now articulated as a genuine concern, sitting underneath multiple ideological positions and binding them in a way the original model does not anticipate. In Chapter 5, this pattern is identified as one of the strongest candidates for an additional structural feature of the contemporary ideological field.

4.7 Synthesis: Configurations and Cleavages

The six patterns documented above, together with the additional patterns recorded in Appendix F, do not occupy the analytical field independently of one another. Instead, they cluster, reinforce, and contradict. This closing section names the configurations and cleavages that the patterns generate.

Three main configurations regularly surface. The first is the qualified Techtopian-Work Machine configuration, exemplified most clearly by Alexander Ietan, John Stringer, Markus, Matteo, and Nadia. In this configuration, respondents articulate strong Techtopian belief in the historical significance of technology and Work Machine endorsement of its productive applications. Meanwhile, they reserve qualifications around surveillance, provenance, and (variably) the labour transition. The closing optimism is robust, the home is broadly accepted as a legitimate site of deployment and the human form factor is treated with curiosity rather than resistance. The second is the ambivalent configuration, demonstrated by Tomáš, Elena, Charlotte, Stefan,

Sofia, and Liliana. In this configuration, respondents acknowledge the productive potential of technology while developing sustained critical reflection on its existential consequences. The home is largely defended, the work-purpose distinction is articulated fully, cognitive atrophy is foregrounded, the closing optimism is reached after substantive work. The third is the resistant-projective configuration, exemplified by Henrik, Petar, and (in a mild form) Isabel. In this configuration, respondents articulate genuine moral and emotional opposition to the technology while remaining unable, even in their own self-description, to detach from it. They resist but that very resistance is held in place against a broader sense of inevitability. In contemporary society it becomes unavoidable to engage with technology. This includes humanoid robots in the future.

The cleavages cut across these configurations along several axes. The strongest is the home cleavage (Appendix F.1): Should the domestic interior be defended or industrialised? The second is the attachment cleavage (Appendix F.4): Is human-robot bonding an extension of capacity or a symptom of pathology? The third is the velocity cleavage (4.4): Is the pace of change exhilarating or anxiety inducing? The fourth is the UBI cleavage (Appendix F.3): Can income alone substitute for the meaning currently supplied by work? These cleavages are not derivative of the configurations. Respondents within the same configuration can fall on opposite sides of any of them.

The pattern that runs underneath all of the above is qualified optimism with which the chapter opened. It functions as a terminus the corpus collectively gravitates toward. It is the closing position toward which substantive ambivalence, sustained critique, and even genuine resistance ultimately bend. Chapter 5 takes up the question of what this terminus, and the patterns that lead to and through it, imply for the adequacy of Kozinets' (2008) ideological field model in the present technological moment.

5 Discussion

This chapter connects the empirical findings of Chapter 4 with the theoretical framework set out in Chapter 2. Its task is to evaluate the adequacy of Kozinets' (2008) ideological field model in the light of the present corpus. Where the data allows it, the chapter proposes modifications and extensions that a contemporary mapping of the ideological technology consumption field appears to require. The chapter is organised around the three theoretical questions posed in Chapter 1. How do the four ideological elements manifest in consumer narratives about humanoid robots? What patterns of ideological movement, contradiction, and co-articulation emerge from the corpus? Does the original model retain its analytical validity under contemporary conditions?

The argument develops in four movements. The first considers the elements that the data affirm. The second considers the points at which the data challenge the model. The third considers the extensions the data require. The fourth integrates these observations into a proposed revised model that retains Kozinets' underlying architecture while incorporating the contributions identified by this thesis. A short closing section reflects on the broader theoretical implications for consumer culture theory in an accelerating technological moment.

5.1 What the Data Affirm

Fundamentally, the data affirm that the four ideological elements appear to remain analytically operative. The Tectopian, Green Luddite, Work Machine, and Techspressive positions identified by Kozinets (2008) are recognisable and recurrent in the present corpus. Eighteen years after the original study, the four elements still organise the ideological field through which respondents articulate their relationship to technology. This is a substantive finding in its own right. It indicates that the underlying cultural categories Kozinets identified appear to have analytical durability beyond the particular technological context in which they were originally developed.

The data also affirm Kozinets' central methodological observation that consumers are fluid across the elements rather than fixed within them. Respondents in the present corpus move between Tectopian endorsement, Work Machine concern, Green Luddite resistance, and Techspressive engagement within single interviews and often within single conversational segments. Alexander Ietan articulates Tectopian commitment throughout most of his interview, yet closes with a Green-Luddite-inflected escape fantasy about retreating to Alaska. Sofia expresses Work Machine concern about teaching displacement, then articulates a Techspressive curiosity about classroom

applications. The methodological commitment of the original study, that the ideological field is dynamically inhabited rather than statically partitioned, appears in the present corpus to be confirmed. Furthermore, the dynamic fluidity may be intensified in the contemporary moment. The proliferation of available positions and the velocity of technological change combined with the rhetorical pressure to maintain a defensible stance under conditions of rapid arrival all appear to amplify the movement Kozinets documented.

The model's account of inner contradictions within each element is also affirmed, particularly with respect to the Tectopian and Work Machine elements. The Tectopian inner contradiction Kozinets identifies (between the morally serious tone of progress discourse and the historical record of technological disasters) is articulated across the corpus. This is often done through the same contemporary catastrophic templates: cybersecurity collapse, surveillance state, weaponisation, the Terminator scenario. The Work Machine inner contradiction Kozinets identifies (between productivity and the shadow side of enslavement and exploitation) is, as Appendix F.3 documents, no marginal qualification but the centre of the Work Machine articulation in this corpus. Respondents who endorse Work Machine reasoning lead with the contradiction rather than reaching it as an afterthought. The model's identification of these contradictions is correct. The data suggest, however, that the contradictions have moved from the periphery to the centre of the respondents' articulations.

A further affirmation concerns the semantic relations among the elements. Kozinets specifies four contraries (morality, individualism, standards, indulgence) and two complementarities (reason, emotion). The data confirm the operative force of the morality contrary (the moral seriousness of Tectopian and Green Luddite positions and their fundamental disagreement about whether technology consumption is good or bad for society) through the home cleavage analysed in Appendix F.1. They confirm the indulgence contrary (the difference between Work Machine instrumentalism and Techspressive pleasure) through the recurring distinction respondents make between the robot as functional tool (acceptable to Work Machine respondents) and the robot as personal extension (acceptable to Techspressive respondents). The complementarity of reason between Tectopian and Work Machine is visible throughout the qualified Tectopian-Work Machine configuration identified in section 4.7, in which the same respondents articulate both elements within a single consistent narrative. The complementarity of emotion between Green Luddite and Techspressive is more diffusely evidenced in this corpus, but is recognisable in both the resistant-

projective and the attached-pro-attachment articulations. The relational structure Kozinets specifies is therefore broadly affirmed.

5.2 What the Data Challenge

The data also surface several patterns that strain components of the original model. Three challenges deserve specific attention.

The first challenge concerns the Techtopian element's mode of operation. The qualified optimism paradox documented in section 4.1 suggests that the Techtopian closure observed across the corpus is often performative rather than substantive. Respondents declare themselves optimistic in the closing register of their interviews even when the content of the preceding conversation has been substantially critical. The Techtopian element appears to function less as a stable ideological position than as a culturally enforced terminus of permitted speech about technology. Kozinets (2008) treats the four elements as analytically equivalent positions in the ideological field. However, the present data suggest that the Techtopian element may occupy a structurally privileged location: it may be the position into which respondents arrive when no other resolution is available. This privilege appears to be cultural. It reflects the present cultural conditions of technology discourse and it demands explicit incorporation into the model's account of how the field is engaged.

The second challenge concerns the structural location of inner contradictions. As mentioned in the previous section, the inner contradictions of both the Work Machine and Techtopian elements have moved from the periphery to the centre of how respondents articulate those elements. The contemporary Work Machine respondent does not endorse productivity and considers meaning an afterthought. The meaning question is raised quickly, often before completing the productivity case. The contemporary Techtopian respondent does not enjoy the technological sublime and then concede, with embarrassment, the historical record of disaster. Disasters are named early and optimism is expressed despite them. This relocation of contradiction is not anticipated by the 2008 model, which treats contradictions as structurally significant but analytically secondary features of each element. The present corpus suggests that under conditions of rapid and visible transformation, the contradictions become the primary phenomenological texture of the ideological articulation, with the supreme good occupying a less central position than the model assumes. This challenges not the model's catalogue of contradictions (which the data confirm) but its assumption about how the contradictions operate within respondents' active articulations.

The third challenge concerns the Green Luddite element's historical baggage. Kozinets (2008) identifies the Green Luddite element as the only position within the field from which to oppose technology consumption, and notes that it carries stigmatic associations with "unfashionable, austere, or violent historical actors". The present corpus contains several respondents (Tomáš, Isabel, Charlotte, Sofia, Stefan) who articulate substantively Green Luddite content, defending the human, the home, and embodied human capability against technological encroachment. They do so without identifying with, or in many cases being aware of the historical Green Luddite tradition. The contemporary Green Luddite articulation is less an inheritance from a recognisable counter-cultural lineage than an emergent ideological position that respondents are constructing in response to specific contemporary provocations. The position retains what Kozinets identified as its distinctive function (providing the only available opposition) but appears to be detaching from its historical iconography. This is a challenge not to the existence of the element but to the model's account of what feeds it culturally. In the present corpus, the element is fed by reflexive observation of contemporary technological dynamics rather than by identification with a counter-cultural heritage. Hence, the implication is that the Green Luddite element may be considerably more durable than its stigmatic associations would suggest. It is being continuously regenerated from contemporary experience rather than depending on inherited cultural material.

5.3 What the Data Extend

This present study proposes five extensions of the model.

5.3.1 The Surveillance Substrate

The first extension concerns the role of surveillance in the contemporary ideological field. As section 4.2 documents, surveillance appears as to be the most cross-cutting concern in the corpus. It is articulated by respondents from across the ideological spectrum and it converges on a common analytical core. The 2008 model does not foreground surveillance as a structural feature of the ideological field. The framework's treatment of privacy is incidental rather than constitutive. The contemporary data demand a more structural account.

The present study proposes that surveillance now operates as a substrate of the ideological field. A substrate is a structural condition that is not internal to any single element, instead it conditions the operation of all four. The Techtopian respondent, affirms the progressive trajectory of the technology, but cannot do so without paying a price in terms of surveillance concerns. The Green Luddite respondent grounds

resistance in surveillance concerns. The Work Machine respondent locates the threat partly in the surveillance dimension of automated systems. The Techspressive respondent withdraws from products whose surveillance practices are perceived as objectionable. Surveillance is therefore not the property of one ideological position. It is the substrate on which all four positions are articulated.

The extension is significant because it implies that the four elements no longer occupy the field independently. Each is now anchored to a shared substrate. That substrate constrains its articulation, supplies a vocabulary of concern that is recognisable across respondents irrespective of ideological position, and generates a degree of analytical convergence the 2008 model does not predict. The surveillance substrate is, in this account, a structural condition of contemporary ideological speech about technology, not an inner contradiction of any specific element.

The surveillance concern rarely appears as a freestanding position. It appears instead as the qualifying clause attached to a position, the recurring “but” and “if” through which respondents from Lukas to Charlotte to Maya articulate their substantive stances. This signature also clarifies the analytical difference between a substrate and an inner contradiction, which the revised model in section 5.4 treats as different kinds of structural feature. A contradiction has paradox form: it arises from the believer’s own commitments and destabilises the position from within. A substrate has conditional form. Emergent, because of the historical moment, it pressures every position and raises the cost of articulation without redirecting it.

5.3.2 The Provenance Variable

The second extension relates to the role of provenance. Both geopolitical and corporate origin plays a fundamental factor in the trust calculation that respondents bring to humanoid robotics. As section 4.3 documents, the question of who makes the robot operates as a structural input to the respondent's overall ideological positioning. Respondents who are otherwise Techtopian may withdraw their endorsement when the technology is associated with provenance they distrust. Respondents who are otherwise Green Luddite qualify their resistance when European or otherwise trusted provenance is invoked. The 2008 model does not specify provenance as a structural feature of the ideological field.

The present study proposes that the contemporary ideological field requires a provenance variable. That is, a structural attribute of the technology, that conditions respondents position toward it. The variable operates along two axes. The first is geopolitical, organising respondents' positions toward Chinese, American, and

European provenance respectively. The second is corporate, organising positions toward specific firms (Tesla, Anthropic, OpenAI, the Chinese manufacturers, etc.) on the basis of perceived value alignment. Provenance does not replace the four elements. Instead, it modifies them. Kozinets' Techtopianism in 2008 could afford to treat the producer of the technology as a neutral background variable. Eighteen years later, in a corpus where Bruno articulates differential trust between Anthropic and OpenAI as a determinant of his substantive stance, that neutrality no longer holds.

Importantly, the provenance substrate is valence-neutral. It does not uniformly burden the positions it conditions. For Matteo, American provenance operates as a reinforcement of an already favourable stance, whereas for Liliana only German or European provenance renders the device acceptable at all. The two respondents would not agree on which origin to trust, yet both treat origin as a gate the evaluation must pass. Even the minority position documented in section 4.3, in which Martin Reisenbichler and Charlotte foreground security architecture over national origin, confirms the substrate. Reasoned indifference is itself a position on provenance, articulated within a frame the question has made unavoidable. What is universal is not the answer but the question. In the present corpus no respondent finds the provenance question unreasonable, and every respondent has pre-formed opinions to some extent. This is precisely the empirical signature of the substrate.

On top of that, the provenance variable is analytically related to the surveillance substrate but is not reducible to it. Respondents distinguish between the fact of surveillance, which surveillance-substrate analysis addresses, and the destination of surveilled data, which is addressed by provenance. The institutional context of production materially conditions the trustworthiness of the device, and that context is no longer ignorable in the way the 2008 model implicitly assumes.

The author acknowledges that provenance concern is not in itself historically novel. Country-of-origin effects on product evaluation have been documented since the 1960s (Schooler, 1965), and the literature was already mature enough to warrant review four decades ago (Bilkey & Nes, 1982). What the present corpus suggests, however, is a transformation in the mechanism. The classical country-of-origin effect operated at the moment of purchase: the consumer judged the origin of the product at the shelf, and the product's connection to its origin effectively closed once it entered the home. The humanoid robot, by contrast, remains tethered to its origin after purchase. It is updated from there, governed from there, and, in the fear Felix articulates, potentially deactivated from there. Provenance has accordingly shifted from a purchase-moment heuristic to a condition of ownership. The surveillance

substrate supplies the mechanism of the tether: the origin matters to respondents in large part because of what the device perceives. Lastly, a methodological qualification is warranted. Part of the provenance material was elicited by a direct interview question about origin, and the substrate claim accordingly rests most firmly on the respondents, among them Alexander Ietan, Felix, and Liliana, for whom the concern surfaced spontaneously.

5.3.3 The Velocity Dimension

The third extension concerns the temporal modality of the ideological field. As section 4.4 documents, the pace of technological change is a structuring feature of respondent narratives about humanoid robotics. Respondents do not only have positions about what the technology is and what it should be. They also have positions about how fast it is arriving and how that pace should be evaluated. The pace is therefore both an object of analysis and a condition that shapes the analysis itself. The 2008 model is largely silent on this dimension.

This thesis proposes that the contemporary field requires a velocity dimension. This is an analytic axis along which respondents are positioned according to their orientation toward the rate of technological change. Some respondents occupy what could be termed a velocity-positive position. They experience the pace as exhilarating and identify with the leading edge of change. They treat the speed itself as a defining feature of the moment they want to inhabit. The other end of the spectrum is the velocity-anxious position. These respondents experience the pace as overwhelming and see themselves as struggling to keep up. They treat the speed as a threat to the conditions needed for careful, considered engagement with the technology. The velocity dimension cross-cuts the four substantive elements. A Techtopian respondent can be velocity-positive (John Stringer) or velocity-anxious (Stefan, who endorses the trajectory but is anxious about its pace). A Green Luddite respondent can fall on either side in the same way.

The velocity dimension matters because it produces a kind of ideological cleavage that the content of the four elements cannot capture. Two respondents can hold identical views about the productive applications of humanoid robotics, the need for regulation, and where corporate responsibility should sit. They can still stand in completely different affective relationships to the same field. One experiences the present moment as a privilege of historical timing. The other experiences it as a chronic source of anxiety. The 2008 model does not anticipate this cleavage, because the rate of technological change was not as pronounced then as it is now.

Two refinements of the velocity dimension warrant explicit statement. The first concerns its relationship to the elements. There is a clear ideological affinity at work. Techtopian commitments incline a respondent toward a velocity-positive position, as John Stringer and Markus illustrate. Resistant commitments incline toward a velocity-anxious one, as Tomáš illustrates. The affinity, however, does not amount to entailment. Felix is broadly progressive in his orientation toward the technology, yet he argues that “we should figure out what we’re doing before just doing it” and describes a coordinated pause as desirable though unrealistic. He endorses the destination while objecting to the speed of travel. The four-element model cannot describe this configuration without contradiction. The reverse case is equally instructive. A Green Luddite is not satisfied by a slower rollout, because the objection is to the technology and not to its rate. A respondent’s velocity position is therefore not derivable from their elemental position. It is this dissociability in the data that warrants treating velocity as a dimension of its own.

The second refinement is that the dimension contains two separable aspects: the perceived rate of change and the respondent’s stance toward that rate. These also come apart in the data. Liliana puts the horizon for domestic humanoid adoption at decades away and so perceives the field as slow rather than accelerating. Lorenz doubts, on resource and energy grounds, that the technology will ever reach truly broad scale. Felix suspects that parts of the industry exaggerate the pace, “throwing in buzzwords without having the actual technology to back them up”, while still advocating deliberation. Henrik feels the acceleration acutely. He responds not by demanding deceleration but by deliberately withdrawing from the information stream. Karin, who keeps the technology at a distance, reports no overwhelm at all. The resistant pole of the corpus therefore scatters across the velocity dimension rather than clustering at its slow-down end. This scattering is stronger evidence for the dimension’s independence than any single deviating case would be. The author also notes that pace-concern is not in itself historically novel. What distinguishes the present field is that velocity has become an openly contested object of consumer discourse, with named positions and public calls for deceleration. The temporal unease of Kozinets’ informants, by contrast, was absorbed into the elements’ inner contradictions and into the model’s account of dynamic fluidity.

A final observation concerns the source of velocity positions. Within this corpus, a respondent’s stake in the field appears to predict their velocity position more reliably than their proximity to the technology. The respondents who embrace acceleration most fully, among them John Stringer, Markus, and Alexander Ietan, are those whose

professional identity, community membership, or competitive position depends on the acceleration narrative being true. The most technically informed respondents, by contrast, tend toward caution. Lukas states this inversion as an explicit in-corpus theory: “the people most cautious about adoption tend to be tech-savvy people who understand the risks”. Emily Kate Genatowski, the most directly involved respondent of all, welcomes the slowness of adoption precisely because it leaves time to regulate. In Kozinets’ (2008) vocabulary, velocity positions appear to be interpellated like any other position in the field. The accelerationist stance recruits those whose subject position is articulated through the technology’s success. The author advances this reading with deliberate caution. Part of the sample was recruited through technology-enthusiast networks, as disclosed in section 3.6. An involvement-breeds-embrace intuition is therefore exactly the kind the counter-intuitive search procedure exists to discipline, and Lukas’ testimony is cited here against it. The stake hypothesis is accordingly advanced as a direction for further research rather than as a finding, and section 6.4 returns to it.

5.3.4 Cognitive Atrophy as a Fifth Inner Contradiction

The fourth extension concerns the structural location of cognitive atrophy in the contemporary field. As section 4.6 documents, a substantial subset of respondents articulate an explicit anxiety about the cognitive consequences of widespread artificial-intelligence and robotic adoption. Some of these respondents are themselves heavy users of the relevant tools. The fear is reflexive rather than externally accusatory. It locates the problem in the user rather than the technology. Kozinets (2008) identifies the inner contradictions of each element but does not foreground cognitive atrophy as one of them.

The present study proposes that cognitive atrophy now operates as a cross-cutting inner contradiction. It runs through multiple elements rather than belonging to any one. Within the Techtopian element, it appears as the worry that the very Progress the position endorses may produce cognitive impairment in those who benefit from it. Within the Work Machine element, it appears as the concern that automation, while productive, may erode the cognitive capacities that make productive work meaningful. Within the Techspressive element, it appears as the suspicion that the pleasures the technology delivers are themselves cognitively diminishing. The contradiction is structural rather than incidental. It appears to be constitutive of how the contemporary field is articulated.

The classification of cognitive atrophy as a contradiction rather than a third substrate warrants explicit justification, since the concern, like the substrates, touches multiple elements. The distinguishing criterion is logical form. A substrate conditions a respondent irrespective of their commitments: the robot collects data whether or not its owner cares. A contradiction, by contrast, requires the believer's own commitments in order to bite, because the erosion of a capacity only wounds a position that values the capacity. Cognitive atrophy has paradox form throughout the corpus: the tool adopted to extend the mind ends up eroding the mind. It is cross-cutting only because several elements share the commitment it bites. The corpus also indicates where it does not bite. Within the Techspressive element the evidence is present but qualified. Maya describes the erosion of interpersonal capacities through technologically mediated comfort. This should be distinguished from the element's classical contradiction of pleasure collapsing into compulsion, which she also voices. Furthermore, no respondent articulates the concern from a nature-grounded Green Luddite position. The classification itself predicts this absence: atrophy is a contradiction of delegation, and the one position that refuses the delegation has nothing to atrophy through it.

The corpus additionally shows that each element answers the shared contradiction with a coping strategy drawn from its own internal logic. Adrien, whose orientation is robustly Techtopian, proposes symbiosis: humanity must grow with the technology rather than against it, up to and including neural integration. Elena and Tomáš, whose registers are pragmatic and deliberative, propose disciplined use: verification, sparing deployment, and the provocation of thought rather than its substitution. Henrik proposes an abstention rule, declining the tools for tasks he already masters. Lorenz proposes analogue reversion, returning to the paper notebook for capacities he wishes to preserve. One contradiction produces four element-consistent management strategies. The pattern demonstrates two things at once. It shows that the contradiction is shared, which the 2008 architecture of element-specific contradictions does not anticipate. It also shows that the elements remain analytically distinct, since the responses to the shared contradiction diverge along elemental lines.

The proposal here is not that cognitive atrophy constitutes a fifth ideological element. It does not generate a distinct nodal point or supreme good of its own. It functions, rather, as an inner contradiction that the contemporary moment has activated across the elements that Kozinets identifies. But the activation is structural and recurring enough that the model should explicitly accommodate it, alongside the inner contradictions Kozinets specifies.

Finally, cognitive atrophy connects to the human-reservation position developed in section 5.3.5 in a structurally revealing way. It is the reservation concern turned inward, and it is the one encroachment against which a boundary strategy fails. Every other reserved territory can be defended at a threshold. The mind, by contrast, is entered by its own owner, voluntarily and daily. The boundary-keeper, in this single case, is the one holding the gate open.

5.3.5 The Human-Reservation Position

The fifth and most theoretically consequential extension concerns the emergence of what the present study terms the human-reservation position. The pattern is documented across multiple findings sections in Chapter 4: the home as ideological boundary (Appendix F.1), the uncanny valley as ideological trigger (Appendix F.2), human-made value as resistance ideology (section 4.5), and elements of the cognitive-atrophy concern (section 4.6). Taken together, these findings appear to surface an ideological position that does not map cleanly onto any of Kozinets' four original elements. It is not Green Luddite in the historical sense. Respondents who articulate it do not oppose technology in general, are not nature-aligned, and do not invoke environmental destruction as their primary frame. It is not Work Machine in either its pessimistic or its qualified-acceptance modes, since the concern is not about labour displacement or productivity. It is not Techtopian, even in inverted form, because there is no underlying commitment to Progress as a supreme good. And it is not Techspressive, in either its enthusiastic or its disengaged forms, because the relationship to the technology is not aesthetic or expressive.

The most natural objection to the proposal is that the position is simply a contemporary intensification of the Green Luddite element, whose supreme good of natural ways of life also covers tradition. To meet that objection, four operational tests are stated here. On each of them, the two positions come apart in the corpus. The first test concerns the warrant: not what is preserved but why it is valuable. Luddite preservation derives value from naturalness and rootedness. Human-reservation derives it from humanness, irrespective of age. Emily Kate Genatowski is the decisive case. She is the most aggressively anti-traditional respondent in the corpus and advocates the industrialisation of domestic space, yet she draws the sharpest boundary in the dataset: replace the work, not the worker. A Luddite reading cannot hold her; the human-reservation reading can. The second test concerns scaling. Luddite resistance scales with the quantity of technology. The human-reservation concern is indifferent to quantity and sensitive only to placement. Sofia endorses robotic deployment in workspaces and refuses it at home, a boundary rather than a gradient. The third test

concerns the square's own relations. If the position were the Luddite corner intensified, it would inherit that corner's contrariety of morality with the Tectopian element, and contraries cannot be held together stably. The corpus shows the opposite. Alexander Ietan, Sofia, and Emily Kate Genatowski each hold Tectopian endorsement and the reservation at the same time, partitioned by domain rather than oscillating. The fourth test concerns temporal direction. Luddite preservation is restorative and faces the past. The human-reservation articulations documented in section 4.5 are, in that section's own framing, confident anticipation rather than nostalgic refusal. They negotiate which domains remain human inside a future the speakers otherwise welcome. The contrast can be compressed into a single line. The Green Luddite fears that technology will degrade human life, whereas the human-reservation position fears that technology will replace the human in it. The author additionally observes that, in the present corpus, ecological framing and human-reservation framing do not appear to co-occur within the same respondents. This is a dissociation an intensification reading would not predict, though it warrants systematic verification against the coded corpus.

The human-reservation position is organised around a supreme good that the four-element model does not specify: the irreducibility of the human. Respondents who articulate this position hold, with varying degrees of explicitness, that there are dimensions of human existence that should not be displaced, replaced, or substituted by humanoid technology. These dimensions are the relational, the embodied, the creative, the inwardly cognitive, and the morally responsible. The position is normative rather than predictive. It is not the claim that humanoid robots cannot perform certain functions. It is the claim that they should not, and that the human dimensions they would displace are dimensions whose preservation is itself a moral and cultural value. The position is found in the most theologically resistant respondents (Tomáš), the most domestically protective (Isabel, Charlotte), the most historically reflective (Lorenz), and the most professionally exposed to AI and robotics integration (Alexander Ietan, in the escape-fantasy register).

The semantic relations between the human-reservation position and the four original elements warrant explicit specification, since the position is best understood relationally rather than substantively. With the Tectopian element, the relation is one of tension over scope. Tectopian endorsement of progress and human-reservation insistence on preserved domains can coexist within a single respondent. They generate productive friction over how far the technology should be permitted to extend. Alexander Ietan's articulation embodies this tension within a single voice: enthusiastic

Techtopian endorsement of robotic integration in his professional domain, paired with a closing fantasy of retreat to Alaska. With the Green Luddite element, the relation is one of partial alignment on different grounds. Both positions resist technological deployment in particular domains. The human-reservation grounds the resistance in the irreducibility of the human, while the Green Luddite grounds it in the integrity of nature or tradition. The two positions sometimes reach the same conclusion by different argumentative routes. With the Work Machine element, the relation is one of tension over what counts as economic value. Work Machine treats productivity and efficiency as supreme goods. Human-reservation treats certain inefficiencies as themselves goods, on the view that the human work of accomplishing the task is part of what is being valued. With the Techspressive element, the relation is one of tension over the nature of pleasure. Techspressive treats technologically optimised pleasure as a good. Human-reservation treats certain pleasures as constitutively human, and therefore valuable precisely because they are not technologically reproducible: handmade objects, human-rendered service, the embodied presence of another human in the same room. The four relations are not contrarities in the strict Greimasian sense. They are tensions that the contemporary moment generates, and that the human-reservation position is, in the present corpus, still in the process of articulating. For this reason the position is best described as a satellite element. It is bound to the field and in sustained tension with all four of its positions, yet it is not generated by the oppositions that produce the square. The square's two axes are oppositions about technology: whether it is good for society, and what its good consists in. The human-reservation position is anchored to a different nodal point, the human itself. It therefore cannot occupy a corner of a square it is not generated by. This is the analytical claim the orbital relation in Figure 3 represents: the four tensions are described here individually and drawn collectively.

The proposal here is that the human-reservation position constitutes a fifth ideological element that the contemporary technology consumption ideological field appears to require. It has its own supreme good (the irreducibility of the human), its own internal logic (preservation, rather than mastery, progress, autonomy, or aesthetic engagement), and its own inner contradiction. The contradiction is that respondents who articulate this position frequently rely, or expect to rely, on the very technologies they ideologically resist. Sofia, who insists that human relational practices must be guarded, will probably encounter humanoid robotics in the school environment within a decade. Isabel, who categorically refuses the prospect of robotic childcare, lives in a society in which such care is being developed. Lorenz, who endorses the philosophical analysis of Kaczynski, is conducting the interview through a digital platform. The

contradiction does not invalidate the position. It locates the position within the same paradox-management dynamic that Mick and Fournier (1998) identify as constitutive of consumer experience of technology more generally.

In the present corpus, the human-reservation position has the standing of an emergent rather than a fully consolidated ideological articulation. Its respondents do not yet recognise themselves as occupying a shared position. The position is being constructed in real time, in response to specific contemporary provocations, by respondents who in many cases would not previously have articulated anything resembling it. The empirical observation that motivates the proposal is that the position now appears with enough structural consistency, across enough respondents, in relation to enough distinct objects (home, body, relationship, artefact), to warrant analytical recognition. The timing of the emergence is itself explicable. A boundary is discovered only when something attempts to cross it, and no prior consumer technology claimed the human form, the caregiving role, or the domestic interior in the way the humanoid robot does by definition. The position Kozinets had no occasion to observe in 2008 is, in the view of the author, being born in response to precisely this claim.

5.4 A Revised Model

The four affirmations, three challenges, and five extensions developed above can be integrated into a proposed revision of the 2008 model. Figure 3 presents the assembled structure visually, in the layered form previewed in Figure 2. The remainder of this section unpacks its components. The revision keeps Kozinets' underlying architecture entirely unchanged: the four-element ideological field, the semantic relations among the elements, and the dynamic fluidity of consumer movement. To this it adds structural features of three distinct kinds, a satellite element, two substrates, and an orthogonal dimension, together with one cross-cutting contradiction.




Figure 3 The orbital ideological field model of technology consumption in an accelerating tech era. Elements and semantic relations adapted from Kozinets (2008).

The proposed revised model has three modifications relative to the original. Firstly, the field now contains five elements rather than four. The Tectopian, Green Luddite, Work Machine, and Techspressive positions are preserved, and a fifth element, Human-Reservation, is added. The five elements are not analytically equivalent. The Tectopian element retains its culturally privileged status as the dominant rhetorical terminus of permitted speech about technology. The Human-Reservation element occupies the emergent position the corpus surfaces but does not yet fully consolidate. The fifth element stands off the square as a satellite, for the reason specified in section 5.3.5: it is anchored to the human rather than to the technology, and is therefore bound to the field without being generated by the square's oppositions.

Secondly, the field is now structured by two substrates that condition all five elements rather than belonging to any one. The surveillance substrate organises the data-privacy and observational concerns that respondents from across the field share. The provenance substrate organises the geopolitical and corporate trust calculations that respondents from across the field bring to the technology. The substrates are not

themselves ideological positions. They are structural conditions on the articulation of the five positions.

Thirdly, the field is now organised along an additional velocity dimension: an analytical axis that positions respondents according to their relationship to the rate of technological change. The dimension is orthogonal to the five elements. A respondent's position on the velocity dimension is not fixed by their position on the substantive ideological elements. The elements incline velocity positions without determining them, in the manner specified in section 5.3.3.

The additional inner contradiction of cognitive atrophy sits across multiple elements rather than within any one, in the manner specified in section 5.3.4. It is not a structural feature of the field in the way the substrates and the velocity dimension are. It is a recurring contradiction that the contemporary moment has activated across the field.

Read in this layered form, the model operates as an analytical instrument rather than only as a map. Three questions, one per layer, can be put to any consumer narrative in the field. Firstly, where does the speaker stand on the field plane, and which contradiction is showing? Secondly, how do the substrates condition the articulation: as burden, as reinforcement, or as priced-in indifference? Thirdly, what rate of change does the speaker perceive, and what is their stance toward it? A single composite utterance illustrates the procedure. Imagine a respondent who says that robots in elder care would be fantastic, but that she would not trust one built by a particular firm, that it would be recording everything in her flat, that the whole development is moving too fast, and that a robot should in any case never be the one holding her mother's hand. She has produced five clauses belonging to four different layers of the model: a Techtopian endorsement, a provenance qualification, a surveillance qualification, a velocity stance, and a human-reservation boundary. The four-element model reads such an utterance as confusion or oscillation. The revised model reads it as a well-formed sentence of the contemporary field.

Table 2 summarises the five elements along identical dimensions to permit direct comparison.

Table 2 The five elements of the revised model along identical dimensions.

Element	Supreme good	Internal logic	Inner contradiction	Characteristic objects in this corpus
Techtopian	Progress	Advancement of society through technological development	Utopian promise against the record of failed delivery; in this corpus, cognitive atrophy	Elder care, healthcare, infrastructure
Green Luddite	Nature and natural ways of life	Principled opposition to technological mediation	Unavoidable reliance on the technology being resisted	Environment, craft, the pre-digital everyday
Work Machine	Economic growth and efficiency	Instrumental evaluation of technology as productive input	Productivity shadowed by enslavement and loss of control; cognitive atrophy	Workplace automation, logistics, labour markets
Techspressive	Pleasure and self-expression	Technology as a medium of enjoyment and identity	Pleasure collapsing into compulsion; interpersonal atrophy (qualified)	Gaming, social media, consumer devices
Human-Reservation	The irreducibility of the human	Preservation of domains valued as human	Reliance on the technologies fenced out; the mind as the unfenceable territory	Home, body, relationship, artefact, mind

The qualified optimism paradox identified in section 4.1 is, on the present account, a phenomenological feature of the contemporary ideological field rather than an additional element or contradiction. It registers the way Techtopian closure operates as the default rhetorical terminus of contemporary technology talk, regardless of the substantive content that has preceded it. The implication for analysis is that researchers working with this field should distinguish between substantive and performative Techtopianism. They should not treat the closing optimism as an unmediated indicator of underlying ideological position.

The revised model retains the methodological commitments of the original. It is constructed from a small purposive sample, analysed through a hermeneutic procedure, and offered as a heuristic mapping of the field rather than as a predictive theory of consumer behaviour. Its claim to advance the literature lies in the structural features it adds (the fifth element, the two substrates, the velocity dimension, and the relocated inner contradiction) and in the analytical purchase those features provide for interpreting consumer narratives in a moment of accelerating technological transformation.

5.5 Theoretical Implications

The proposed revision has several broader theoretical implications for consumer culture theory and for the study of technology consumption specifically.

The first implication concerns the analytical status of the human itself in the contemporary technological field. Kozinets' (2008) model treats the four ideological elements as positions through which consumers articulate their relationship to technology. It does not foreground the question of what the human is, in relation to the technology, as itself a structurally significant variable. In the present corpus, however, that question appears to be increasingly at stake. Respondents who articulate the Human-Reservation position are not merely taking a stance on humanoid robotics. They are articulating a position on the irreducibility of the human, at a moment when that irreducibility is being technologically contested. The implication for consumer culture theory is that the human itself may need to be theorised as a variable the technological field acts upon, rather than as a fixed background to that field.

The second implication concerns the temporal modality of the ideological field. The velocity dimension surfaced in this study points toward a more general theoretical observation. Existing consumer culture theory has tended to treat the temporal modality of consumption as a relatively stable background condition, against which substantive positions are articulated. The present data suggest that under conditions

of accelerated technological change, the temporal modality itself becomes a structuring feature of the ideological field. Consumer culture theory may need to develop more systematic tools for theorising velocity, acceleration, and the temporal organisation of ideological articulation.

The third implication concerns the transnational structure of the technological field. The provenance substrate identified in this corpus suggests that contemporary consumer narratives about technology are inescapably structured by geopolitical and corporate-ideological considerations. The more nationally bounded analytical frameworks of earlier consumer culture theory did not foreground these. The implication is that future studies of technology consumption ideology may need to incorporate the geopolitical and corporate origins of the technology under analysis as a structural rather than a contextual variable. The meaning consumers attach to technologies is now substantially conditioned by the institutional ecosystems within which those technologies are produced.

The fourth implication concerns the reflexivity of contemporary ideological articulation. The robot-as-social-mirror pattern identified in Appendix F.6 suggests that contemporary consumers are, with surprising frequency, capable of recognising the projective character of their own articulations and of building that recognition into the articulations themselves. This is at least true of the educated, urban, technologically engaged populations represented in this corpus. It is a meta-level capacity the 2008 model did not have particular occasion to theorise. Contemporary studies of ideological articulation may need to take it more seriously. Consumers are not merely producing ideology. They are increasingly producing ideology about ideology.

The fifth and broadest implication concerns the adequacy of inherited analytical frameworks for a moment in which the technologies under study are themselves transforming the conditions of their own analysis. The humanoid robot discussed in this thesis did not exist in a generally available form when Kozinets developed his 2008 framework. The artificial intelligence that powers contemporary humanoid robots has emerged in commercially significant form only in the last decade. The geopolitical bloc structure that conditions consumer trust in the technology has reorganised substantially in the last five years. Under these conditions, the foundational analytical frameworks of consumer culture theory require not only continuous empirical updating but periodic structural revision. The present study is offered as a contribution to that ongoing revisionary work. The 2008 model remains analytically valuable. It does not, in its original form, suffice for the present moment. The revision proposed here is

intended as one possible direction in which the model can be developed to recover its analytical adequacy under contemporary conditions.

The chapter closes with a methodological observation. The revision proposed here is grounded in twenty-four interviews conducted in early 2026, in a particular cultural and historical moment, by a single researcher operating within a disclosed positionality. It is offered as a heuristic that warrants further empirical examination, not as a finished structural map of the field. The patterns identified in Chapter 4 are reasonably stable across the corpus. The theoretical extensions proposed in this chapter are the analyst's reading of those patterns, and they remain open to refinement through subsequent empirical and theoretical work. Chapter 6 takes up the broader practical, theoretical, and methodological implications of the study, and identifies the directions in which the present contribution most evidently requires further development.

6 Conclusion

This thesis set out to examine how contemporary consumer narratives about humanoid robots align with, challenge, or extend Kozinets' (2008) ideological field model of technology consumption. The empirical material consisted of twenty-four semi-structured in-depth interviews conducted in early 2026, with respondents drawn from across professional, generational, and national contexts. It was analysed through a hermeneutic procedure that combined deductive engagement with the four-element framework and inductive openness to emergent patterns. Twelve cross-cutting patterns were identified, six structuring the main analysis in Chapter 4 and a further six recorded in Appendix F. Chapter 5 brought those patterns into systematic engagement with the original model. It affirmed the model's underlying architecture while proposing a series of structural revisions: a fifth ideological element, two substrates that condition all elements, an additional velocity dimension, and a relocated inner contradiction. This concluding chapter consolidates the contributions of the study, draws out their practical implications for the actors who will shape the public emergence of humanoid robotics, acknowledges the limitations of the work, and identifies the most productive directions for further research.

6.1 Summary of Contributions

The study makes three principal contributions.

The first contribution is theoretical. The thesis constitutes, in the view of the author, the first systematic empirical test of Kozinets' (2008) ideological field model in relation to humanoid robots, conducted in a documented period of acceleration in the field. The four ideological elements (Techtopian, Green Luddite, Work Machine, and Techspressive) appear to be confirmed as analytically operative in contemporary consumer narratives, even though the technology is categorically different from the personal computing focus of the original study. The dynamic fluidity of consumer movement among the elements, which Kozinets identified as the model's most empirically distinctive contribution, is confirmed and, in several respects, intensified. At the same time, the study identifies five structural extensions the model now requires. The proposed revised model retains Kozinets' underlying architecture while incorporating these contributions, and is offered as a heuristic mapping of the ideological field that warrants further empirical examination.

The second contribution is empirical. The study documents six cross-cutting patterns in the main analysis, with a further six recorded in Appendix F:

1. The qualified optimism paradox
2. Surveillance as master fear
3. Geopolitical and corporate origin as trust proxy
4. Velocity anxiety
5. Human-made value as resistance ideology
6. Cognitive atrophy as emergent anxiety

These patterns constitute, in their own right, a contribution to the empirical literature on consumer responses to humanoid robotics. The existing literature has tended to concentrate on attitude constructs derived from the Technology Acceptance Model and its extensions, and to under-attend to the broader ideological formations through which consumers make sense of the technology. The patterns, together with the configurations and cleavages they produce, are available for use, refinement, and contestation in subsequent empirical work.

The third contribution is methodological. The thesis demonstrates the analytical productivity of a hermeneutic part-to-whole interpretation of in-depth interview material, in a study whose theoretical purpose is the simultaneous testing and possible extension of an existing model. The five-cycle analytical procedure described in Chapter 3 (intratextual reading, deductive coding, inductive open coding, intertextual synthesis, and theoretical reflection) is offered as a transferable design for studies of comparable theoretical ambition. The explicit balancing of deductive and inductive cycles is identified as the methodological precondition for studies that aim to honour an existing framework while remaining open to its possible insufficiency.

6.2 Practical Implications

The practical implications of the study are addressed to three constituencies: the firms that will design and manufacture humanoid robots; the marketers and communications professionals who will be responsible for their reception by consumers; and the policymakers who will be tasked with regulating their integration into society.

For firms, the most consequential implication concerns the structural significance of the surveillance and provenance substrates identified in Chapter 5. The data show that respondents from across the ideological spectrum converge on two principal substantive concerns: data privacy, and the institutional accountability of the manufacturer. Some firms may treat these concerns as marketing peripherals, to be addressed through the standard apparatus of privacy policies and brand reputation campaigns. Such firms appear likely to underestimate how structurally located the concerns are in consumer narratives. Several factors are not merely contextual

conditions on the marketing of the device. They are constitutive features of the device as consumers will encounter it: data architecture, encryption, the geographic location of data storage, the specific corporate values communicated by the firm's leadership, and the firm's observable consistency between stated commitments and actual conduct. Firms whose technology stack, corporate identity, and value system align with consumer expectations on the surveillance and provenance substrates will hold a structural competitive advantage. Firms whose offerings do not align, however technically capable, will not. The implication is particularly significant for two groups. The first is European firms, several of whose potential consumers explicitly framed European provenance as a positive trust signal. The second is firms whose corporate identities have been built around publicly articulated value commitments. The differential trust Bruno accords to Anthropic relative to OpenAI is a precise and prospectively generalisable instance.

A related implication for firms concerns the structural location of the home. The data show that the domestic interior is the sharpest cleavage in contemporary consumer narratives about humanoid robotics. Respondents who endorse the technology elsewhere often refuse it for the home. Firms whose go-to-market strategies depend on rapid domestic adoption are likely to encounter substantially more resistance than the workplace and industrial deployment trajectories suggest. On the data of this corpus, the more promising domestic-adoption pathway runs through specific high-value use cases that respondents are prepared to grant: eldercare, accessibility support for people with disabilities, and specific household chores in the form of contained-task tools that can be turned off and stored. It does not run through the generalist humanoid presence that several respondents explicitly resist. Firms designing for domestic deployment may benefit from Emily Kate Genatowski's reframing. The appropriate strategy may be less the introduction of the generalist humanoid into the unmodified home, and more the industrialisation of the domestic space around the robot's capacities. Earlier waves of domestic technology, such as electricity, gas, and the appliance, were accommodated in just this way, through reconfiguration of the home itself.

For marketers and communications professionals, the most consequential implication concerns the qualified optimism paradox documented in section 4.1. The data show that contemporary consumer narratives about humanoid robotics terminate, with overwhelming consistency, in optimistic closure, even when the substantive content of those narratives has been substantially critical. The implication is twofold. On the one hand, the closing optimism represents a cultural disposition that marketers can

address. On the other hand, it masks the substantive ambivalence that has preceded it. Consumer research that relies on closing affective measures appears to risk systematically overstating the public's ideological readiness for the technology. The strategic implication is that communications strategies should address the substantive concerns surfaced in the corpus (surveillance, provenance, cognitive atrophy, the work-purpose distinction, the home) rather than the performative optimism the corpus also generates. The latter is the audience's closing posture. The former is the audience's actual analytical territory.

A second implication concerns the human-reservation position identified as the proposed fifth element. Some marketers may treat consumer attachment to human-made value, human-rendered service, and human-occupied domestic space as residual nostalgia to be educated out. This misreads the structural location of the position in contemporary consumer narratives. The position is not residual. It is emergent and, in this corpus, increasingly articulated. Marketing strategies that explicitly accommodate it are likely to outperform strategies that frame the technology as a totalising substitute. Such strategies will signal where the technology will and will not encroach, preserve identifiable human-rendered alternatives, and frame the technology as a complement to human capability rather than a replacement for it.

For policymakers, the implications are organised around the structural concerns the data surface as already mature in public consciousness. The first is data and surveillance regulation. Respondents across the ideological spectrum want robust regulatory frameworks for the data-collection apparatus that humanoid robotics will introduce into private and public spaces. The implication is not that the public is asking policymakers to obstruct the technology. It is asking them to provide legible structural protections that make adoption psychologically and practically tolerable. Policymakers who treat the regulatory question as a brake on innovation are misreading what the public is actually asking for. The European Union's regulatory direction, in particular, was endorsed across multiple respondents as a comparative advantage Europe should consciously preserve and develop. The endorsement came even from respondents whose own ideological commitments would otherwise have led them to favour minimal intervention.

The second policy implication concerns the labour transition. The data show that respondents are substantially more willing to support reskilling and retraining frameworks than universal basic income. The underlying anthropological objection to UBI, the work-purpose distinction documented in Appendix F.3, is broadly shared even among respondents who are personally well-disposed to the technology. The strategic

implication is that policy responses to large-scale automation may find more public support if they are framed as facilitating labour-market transition than as institutionalising labour-market exit. Policymakers who treat UBI as the natural response to automation should be aware that the public on which the policy is to be enacted does not, in this corpus, share that intuition.

The third policy implication concerns the liability and rights questions that humanoid robotics raises. The data show that respondents are substantially more interested in the immediate question of who is liable when a humanoid robot causes harm than in the more philosophically resonant question of whether robots should have rights. The corpus contains a substantial range of considered positions on the liability question, including detailed elaborations from Emily Kate Genatowski on the inadequacy of the current owner-liable default. Positions on rights are considerably less developed. Most respondents either reject the rights question as a category error or treat it as a distant and speculative matter. The implication for policymakers is that the regulatory work of greatest immediate consequence is the construction of liability frameworks that align with the operational realities of the technology. The rights debate, while important, can reasonably be deferred until the empirical conditions that would activate it are nearer.

6.3 Limitations Revisited

Two further limitations specific to the present study's contribution warrant explicit acknowledgement. The first concerns the Human-Reservation fifth element proposed in Chapter 5. On the data of this corpus, the element is emergent rather than fully consolidated. Respondents do not yet recognise themselves as occupying a shared position. The proposal that it constitutes a fifth element, rather than a recurring articulation across the existing four, is a motivated extension pending confirmation in other product categories and samples. It is not a settled empirical conclusion. The alternative reading, that the human-reservation articulation is best treated as a contemporary intensification of the Green Luddite element, remains worth examining. The four tests stated in section 5.3.5, concerning warrant, scaling, relational behaviour, and temporal direction, tell against it in the present corpus, but the proposal should be understood as a hypothesis the corpus motivates and that further work can test. The second limitation concerns the qualified optimism paradox. The pattern documented in section 4.1 appears robust across the corpus. The interpretation offered in Chapter 5, that the optimism functions partly as a culturally enforced rhetorical terminus rather than as substantive belief, depends on inferential reasoning about the relationship between substantive content and closing register.

Further study could probe this more systematically. Studies that varied the closing question, that compared closing self-reports against measures of substantive position, or that examined the qualified optimism phenomenon across cultural contexts with different conventions of interview closure, would refine the interpretation the present study advances.

6.4 Directions for Further Research

The study suggests several directions for further empirical and theoretical work.

The first direction is cross-cultural extension. The cultural composition of the present sample, while heterogeneous, is dominantly European, with smaller Hong Kong, US, and Eastern European cohorts. The proposed revised model, particularly the human-reservation element, the surveillance substrate, and the velocity dimension, would benefit from systematic testing in cultural contexts whose relationship to technology, the home, and the human is structurally different. Several contexts would be especially clarifying: East Asian contexts where humanoid deployment is already more advanced; South American and African contexts where the geopolitical landscape of provenance is differently configured; and religious-cultural contexts in which the human-reservation position may take distinctive forms. Such studies would substantially clarify which elements of the proposed model are general and which are culturally specific.

A further direction concerns the velocity dimension. The observation in section 5.3.3, that a respondent's stake in the field appears to predict their velocity position more reliably than their exposure to the technology, warrants systematic examination. This would ideally be done in samples whose material and identity stakes in the acceleration narrative vary by design. Such work would also discipline the positionality concern noted there, since the present sample's partial recruitment through technology-enthusiast networks limits how far the stake reading can be pressed on this corpus alone.

The second direction is longitudinal extension. The present study captures the ideological field at a particular moment of rapid public emergence. The natural successor study would re-engage the same field at intervals, for instance after three, five, and ten years. It would document how respondents' articulations evolve as the technology moves from anticipation to encounter, from speculation to deployment, and from the abstraction of media coverage to the concreteness of personal experience. Studies that re-interviewed the same respondents over time would illuminate the individual trajectories of ideological movement that the present

synchronic study can only infer. Studies that introduced new respondents at each interval would illuminate the generational dynamics that the gradient pattern documented in Appendix F.5 anticipates but does not directly test.

The third is methodological extension. The present study used in-depth interviews without pre-stimulus material, on the considered judgement that the spontaneity of respondents' existing mental models was the appropriate object of analysis. Studies that introduced controlled stimulus conditions, such as encounters with humanoid prototypes, exposure to specific media representations, or structured projective tasks, would surface ideological articulations the present design could not. Studies that complemented in-depth interviews with ethnographic observation in contexts where humanoid robots are already present (Emily Kate Genatowski's domestic environment, Alexander Ietan's workplace, the East Asian service-industry contexts that several respondents invoked) would test the present study's findings against situations of actual rather than anticipated encounter.

The fourth is theoretical extension. The proposed revised model is offered as a heuristic that warrants further development. Four directions are particularly productive. The first is the formal articulation of the semantic relations between the Human-Reservation element and the four original elements. The second is the specification of the conditions under which the surveillance and provenance substrates intensify, weaken, or transform. The third is the analytical relationship between the velocity dimension and the dynamic fluidity of consumer movement that Kozinets identifies as the model's central empirical signature. The fourth is the integration of the proposed revisions with adjacent frameworks, including Mick and Fournier's (1998) technology paradoxes and Thompson's (2004) account of marketplace mythologies. Each of these adjacent frameworks intersects the present study's contribution in ways the thesis has only begun to identify.

The fifth is applied extension. The practical implications developed in section 6.2 are addressed to firms, marketers, and policymakers, but they are advanced at the level of strategic principle rather than operational specification. Applied studies that translated the present findings into specific design briefs, communications strategies, and regulatory frameworks would test the analytical purchase of the proposed revised model in concrete decision-making contexts. Such studies would also generate the additional empirical material necessary to refine the model's applied utility.

6.5 Closing Reflection

The thesis began from an observation about a temporal mismatch. The present moment introduces a technology, the humanoid robot, whose ideological provocations appear categorically more concentrated than those of the technologies the original framework was built to analyse. The empirical material collected and analysed in the course of the study confirms this observation in both directions. The original framework retains its analytical validity eighteen years on. At the same time, the contemporary technology substantially exceeds the framework's capacity to fully render the consumer narratives it generates.

The contribution this thesis offers is intended in the spirit in which Kozinets (2008) himself offered the original framework. It is a heuristic mapping of an ideological field that respondents are actively constructing in the moment of being studied, and that the analyst can hope to render legibly enough to be useful to those who will encounter the field in their own work. The revised model proposed in Chapter 5 is an attempt to do for the technological moment of the mid-2020s what the original model did for the technological moment of the mid-2000s. If it succeeds in its analytical work, it will in turn prove inadequate to the technological moment of the mid-2030s, at which point a comparable revision will be necessary. This is not a defect of the framework. It is the proper condition of consumer culture theory under accelerating technological change. The frameworks the discipline inherits warrant continuous renewal, and that renewal is best conducted through the patient hermeneutic engagement with consumer narratives that the discipline's formative work pioneered.

The respondents to whom this thesis is most fundamentally indebted are the twenty-five individuals who, between February and March 2026, gave between thirty and ninety minutes of their lives to a graduate student's questions. Many of them had only just begun to think about the technology, and few had directly encountered it. The richness, sophistication, and ambivalence of their articulations are the substantive content of every analytical claim this thesis advances. The contemporary consumer of technology is not the passive recipient of corporate framings or the manipulated subject of marketing operations. They are active producers of meaning, constructing the significance of new artefacts as they encounter them. The respondents in this study did so with a thoughtfulness and self-awareness that consumer culture theory has long taken seriously, and this thesis is offered as an attempt to do the same.

6.6 A Theoretical Coda: The Interview as Hope-Event

The preceding section closed on the proposition that the contemporary consumer of humanoid robotics is an active producer of meaning, constructing positions in real time in relation to artefacts whose significance is still being determined. The fieldwork conducted in the course of this thesis makes that proposition concrete in a way the recorded transcripts alone cannot fully convey. The aim of this final section is to offer a phenomenological observation about the structure of the interview encounter itself. The observation suggests that the empirical material presented in this thesis is not simply the record of pre-formed ideological positions, but the trace of a hope-event whose terms the philosopher Gabriel Marcel sketched in a different setting almost a century ago.

Three features of the corpus motivate this reading.

The first is that a substantial subset of respondents explicitly marked the interview as the moment at which they had, for the first time, thought their position through. A primary school teacher, before formulating her view, observed: "I've never really thought much about humanoid robots." A respondent confronted with the question of robot rights paused before answering: "Wow, I never thought about that. It's a really tough question." A third, invited to consider robotic care for the elderly, replied: "I've never really thought about it before, but yes. Of course it would be nicer if there were actual humans taking care of them, but I think it's better to have a robot than no one taking care of them." In each instance, the declared absence of prior thought was followed not by silence, but by the production, in real time, of a substantively reasoned position. This observation is consistent with what Kvale and Brinkmann (2009) describe as the latent-cognition-surfacing function of in-depth interviewing, but it appears to be more than that. Respondents were not retrieving stored views. They were constituting their views in the act of speaking them, under the specific pressure of a question they had not previously felt called upon to answer. The interview did not record their position. Instead, it generated it.

The second feature is a structural distinction that several respondents drew between their dispositional stance toward the future and the epistemic situation in which they nonetheless held that stance. The clearest example came from a senior professional with extensive industry exposure: "I'm an optimistic person by nature, but I'm also aware that optimism or pessimism doesn't change what's actually happening. The honest answer is: I don't know, and the uncertainty itself is the most accurate thing I can say." The respondent distinguished, in a single utterance, between a personal

disposition (optimism) and the calculative situation in which that disposition was held (radical uncertainty). A second respondent ended his interview on a similar note: "Overall optimistic, with a lot of pessimism sprinkled in. The future has huge potential, but humanity is just a bit stupid sometimes. That said, we have been through so much trouble and recovered from it. So let us hope that whatever is going on right now is just a phase." The grammatical shift in the final sentence, from prediction to invocation, is the relevant feature: let us hope. A third respondent closed in the same register: "Every day I think, and I hope, let us see." These statements are not optimistic forecasts. They appear to be commitments to a posture that persists when forecasting has reached its limit. The structural location is precisely the one Marcel (1962, originally 1944) reserved for what he called *esperance*, hope as distinct from *optimisme*: a non-calculative posture sustained in the absence of supporting reasons, appearing where calculation cannot. In this light, the qualified optimism paradox identified in section 4.1 admits of a phenomenological re-description. What appeared as optimism qualified by anxiety may in fact be hope mislabelled as optimism. The qualification, far from being a defect of the optimism, may be its identifying signature.

The third feature is grammatical. Marcel's account of hope is irreducibly relational and irreducibly plural. The structure he gives it has three terms: a subject who hopes, a Thou through whom hope is held, and a community for which hope is offered. In his compact formulation: *J'espere en Toi pour nous*. I hope in Thee for us. The pronoun matters. Across the corpus, when respondents moved from describing humanoid robots to expressing their stance toward the future those robots imply, an almost universal grammatical shift occurred from the first-person singular to the first-person plural. "We have to think seriously, as humans and as a society, about what we do with these people." "We will see what happens. We have to be happy somehow, we have to go on." "We have to grow with it. Not against it. The goal of developing it is to help humanity, and if we want it to do that, we have to evolve alongside it." "We have been through so much trouble and recovered from it." In a setting where respondents had every license to speak only for themselves, they instead spoke for a collective subject that included them but exceeded them. The hope was for an us. Several respondents brought the collective into focus through reference to children and future generations. One father commented: "My kids are nine and eleven. They will grow up with full exposure to robots in daily work and daily life. And I hope the robots will take care of us when we are old." Another, considering the disruption ahead: "You wish for the best, especially once you have a child." A third: "It is for the future generation to decide how to handle all of this responsibly." The respondents

were not asking for personal benefit. They were asking that a community to which they belonged would come through.

These three features, articulation as event, the dispositional and epistemic gap, and the collective grammar, converge on the post-interview phenomenon described in Appendix F.6.1. Respondents, having articulated their fears under the structuring asymmetry of the interview, almost universally turned at the close of the recording and asked the interviewer for his own view. The phenomenon was initially read in this thesis as a search for interpretive authority under conditions of cultural unsettlement, and that reading remains correct as far as it goes. Marcel's framework, however, permits a more precise statement. The three terms hope requires were almost present throughout the interviews. The first-person subject was present. The collective subject for whose sake hope was held was present. The middle term, the Thou through whom hope is to be offered, was the term in absence. The interviewer, methodologically committed to non-disclosure of his own ideological framework, had functioned during the recorded portion as a suspended Thou: present as a listener, withheld as an interlocutor. The instant the recording ended, that suspension was released. Respondents reached for the second term they had lacked. Perhaps they were not asking the interviewer for information. They were asking him to be the Thou through whom the hope-structure their utterances had been constructing could finally be completed.

Two consequences follow.

The first is methodological. The post-interview moment is not, on this reading, a residue or by-product of the interview proper. It is the moment at which the interview as a hope-event completes itself. Future qualitative researchers conducting in-depth interviews on technologies that respondents experience as ideologically unsettled may want to take seriously the proposition that the recorded portion is only part of the encounter. The unrecorded moment carries methodological weight of its own. Lincoln and Guba (1985) note that naturalistic inquiry must attend to what the interviewer becomes for the respondent over the course of the encounter. Marcel's framework gives a specific account of one such becoming.

The second consequence matters more, and it is offered here as the final theoretical observation of this thesis. The ideological field this thesis maps is not, on the present evidence, merely unsettled. It is unsettled in a particular structural way. It lacks the middle term required for hope to be performed in common. Respondents have positions. They have collective subjects on whose behalf they wish well. What they do

not have, and what the present cultural moment does not yet supply, is a publicly available Thou: a trusted authority, a recognised interlocutor, an interpretive frame through whom their hopes for the technological future of humanity can be held in common. Bauman (2000) describes liquid modernity as a condition in which solid orientations have dissolved. The present data suggest a more specific phenomenological observation. The dissolution shows up in this particular ideological domain not as despair, but as a vacancy in the second-term position of the hope structure. Respondents are not in despair. They appear to be in the peculiar condition of being able to hope only if someone will be the Thou for them. The thesis interviewer, briefly and unintentionally, found himself in that position. The respondents' near-universal turn to him at the close of each conversation is, in the view of the author, the single most ethnographically significant fact the fieldwork produced. It is also a fact the recorded transcripts alone cannot disclose.

The revised model of the ideological field developed in Chapter 5 should therefore be read with a final qualification. The Human-Reservation element, the surveillance and provenance substrates, the velocity dimension, and the relocated contradiction of cognitive atrophy are all structural features the model now requires in order to render the contemporary moment legible. To these the present coda adds a phenomenological feature that is not strictly part of the model, but is presupposed by the conditions under which the model is being constructed. The ideological field is being articulated by consumers who are reaching, while they articulate it, for the relational structure that would allow their articulation to count as hope rather than as forecasting. The marketing implication, finally, is that firms, communicators, and policymakers who address the humanoid-robotics consumer as if she held a settled position to be persuaded, will miss what the present data show her to be doing. She is performing a hope-structure that is, at the time of writing, only incompletely available to her. She is looking for a Thou. Whether the cultural institutions of the present moment, including the marketing institutions, are capable of being one, is a question the present thesis does not answer. It is, however, the question the data leave standing at the close.

References

- Al Jazeera. (2026, February 17). *Humanoid robots perform advanced martial arts at Chinese New Year gala*. <https://www.aljazeera.com/news/2026/2/17/humanoid-robots-perform-advanced-martial-arts-at-chinese-new-year-gala>
- Althusser, L. (1971). *Lenin and philosophy and other essays* (B. Brewster, Trans.). Monthly Review Press. (Original work published 1968)
- Arnould, E. J., & Thompson, C. J. (2005). Consumer culture theory (CCT): Twenty years of research. *Journal of Consumer Research*, 31(4), 868–882. <https://doi.org/10.1086/426626>
- Arnould, E. J., & Wallendorf, M. (1994). Market-oriented ethnography: Interpretation building and marketing strategy formulation. *Journal of Marketing Research*, 31(4), 484–504. <https://doi.org/10.1177/002224379403100404>
- Bauman, Z. (2000). *Liquid modernity*. Polity Press.
- Belk, R. W. (1988). Possessions and the extended self. *Journal of Consumer Research*, 15(2), 139–168. <https://doi.org/10.1086/209154>
- Belk, R. W. (2017). Qualitative research in advertising. *Journal of Advertising*, 46(1), 36–47. <https://doi.org/10.1080/00913367.2016.1201025>
- Belk, R. W., Wallendorf, M., & Sherry, J. F. (1989). The sacred and the profane in consumer behavior: Theodicy on the Odyssey. *Journal of Consumer Research*, 16(1), 1–38. <https://doi.org/10.1086/209191>
- Bilkey, W. J., & Nes, E. (1982). Country-of-origin effects on product evaluations. *Journal of International Business Studies*, 13(1), 89–100. <https://doi.org/10.1057/palgrave.jibs.8490539>
- Borgmann, A. (1984). *Technology and the character of contemporary life: A philosophical inquiry*. University of Chicago Press.
- Bowen, G. A. (2008). Naturalistic inquiry and the saturation concept: A research note. *Qualitative Research*, 8(1), 137–152. <https://doi.org/10.1177/1468794107085301>

- Burbano, V., Padilla, N., & Meier, S. (2024). Gender differences in preferences for meaning at work. *American Economic Journal: Economic Policy*, 16(3), 61–94. <https://doi.org/10.1257/pol.20220121>
- CNN. (2026, February 18). *Analysis: China's biggest TV event had a clear star: the robot*. <https://www.cnn.com/2026/02/18/china/china-humanoid-robots-new-year-gala-intl-hnk>
- Connell, R. W. (1995). *Masculinities*. Polity Press.
- de Graaf, M. M. A., & Ben Allouch, S. (2013). Exploring influencing variables for the acceptance of social robots. *Robotics and Autonomous Systems*, 61(12), 1476–1486. <https://doi.org/10.1016/j.robot.2013.07.007>
- Ellul, J. (1964). *The technological society* (J. Wilkinson, Trans.). Vintage Books. (Original work published 1954)
- European Commission. (2025). *Special Eurobarometer 554: Artificial intelligence and the future of work*. Directorate-General for Employment, Social Affairs and Inclusion. <https://doi.org/10.2767/8591026>
- Frank, T. (2000). *One market under God: Extreme capitalism, market populism, and the end of economic democracy*. Doubleday.
- Gonalons-Pons, P., & Gangl, M. (2021). Marriage and masculinity: Male-breadwinner culture, unemployment, and separation risk in 29 countries. *American Sociological Review*, 86(3), 465–502. <https://doi.org/10.1177/00031224211012442>
- Greimas, A. J. (1966). *Sémantique structurale: Recherche de méthode*. Larousse.
- Guest, G., Bunce, A., & Johnson, L. (2006). How many interviews are enough? An experiment with data saturation and variability. *Field Methods*, 18(1), 59–82. <https://doi.org/10.1177/1525822X05279903>
- Heerink, M., Kröse, B., Evers, V., & Wielinga, B. (2010). Assessing acceptance of assistive social agent technology by older adults: The Almere model. *International Journal of Social Robotics*, 2(4), 361–375. <https://doi.org/10.1007/s12369-010-0068-5>

Heidegger, M. (1962). *Being and time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row. (Original work published 1927)

Ihde, D. (1990). *Technology and the lifeworld: From garden to earth*. Indiana University Press.

Ivanov, S., & Webster, C. (2019). *Robots, artificial intelligence, and service automation in travel, tourism and hospitality*. Emerald Publishing.

Jahoda, M. (1982). *Employment and unemployment: A social-psychological analysis*. Cambridge University Press.

Jameson, F. (1981). *The political unconscious: Narrative as a socially symbolic act*. Cornell University Press.

Jameson, F. (2005). *Archaeologies of the future: The desire called utopia and other science fictions*. Verso.

Kozinets, R. V. (2008). Technology/ideology: How ideological fields influence consumers' technology narratives. *Journal of Consumer Research*, 34(6), 865–881. <https://doi.org/10.1086/523289>

Kvale, S., & Brinkmann, S. (2009). *InterViews: Learning the craft of qualitative research interviewing* (2nd ed.). Sage.

Laclau, E., & Mouffe, C. (1985). *Hegemony and socialist strategy: Towards a radical democratic politics*. Verso.

Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry*. Sage.

MacDorman, K. F., & Ishiguro, H. (2006). The uncanny advantage of using androids in cognitive and social science research. *Interaction Studies*, 7(3), 297–337. <https://doi.org/10.1075/is.7.3.03mac>

Marcel, G. (1962). *Homo viator: Introduction to a metaphysic of hope* (E. Craufurd, Trans.). Harper Torchbooks. (Original work published 1944)

Markoff, J. (2005). *What the dormouse said: How the sixties counterculture shaped the personal computer industry*. Viking.

McCracken, G. (1988). *The long interview*. Sage.

Mick, D. G., & Fournier, S. (1998). Paradoxes of technology: Consumer cognizance, emotions, and coping strategies. *Journal of Consumer Research*, 25(2), 123–143.
<https://doi.org/10.1086/209531>

Mori, M. (1970). The uncanny valley. *Energy*, 7(4), 33–35.

Munsch, C. L., Rogers, M. L., & Yorks, J. M. (2016, August). *Relative income, psychological well-being, and health: Is breadwinning hazardous or protective?* [Paper presentation]. 111th Annual Meeting of the American Sociological Association, Seattle, WA, United States.

Noble, D. F. (1999). *The religion of technology: The divinity of man and the spirit of invention*. Penguin Books.

Norris, P., & Inglehart, R. (2019). *Cultural backlash: Trump, Brexit, and authoritarian populism*. Cambridge University Press.

Nye, D. E. (1994). *American technological sublime*. MIT Press.

Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). Sage.

Pew Research Center. (2025, April 3). *How the U.S. public and AI experts view artificial intelligence*. <https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/>

Reuters. (2026, February 17). *China's humanoid robots take center stage for Lunar New Year showtime*. <https://www.reuters.com/business/media-telecom/chinas-humanoid-robots-ready-lunar-new-year-showtime-2026-02-16/>

Rushkoff, D. (1996). *Playing the future: How kids' culture can teach us to thrive in an age of chaos*. HarperCollins.

Schooler, R. D. (1965). Product bias in the Central American common market. *Journal of Marketing Research*, 2(4), 394–397.

Saunders, B., Sim, J., Kingstone, T., Baker, S., Waterfield, J., Bartlam, B., Burroughs, H., & Jinks, C. (2018). Saturation in qualitative research: Exploring its conceptualization and operationalization. *Quality & Quantity*, 52(4), 1893–1907.
<https://doi.org/10.1007/s11135-017-0574-8>

- Sharkey, A., & Sharkey, N. (2012). Granny and the robots: Ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1), 27–40. <https://doi.org/10.1007/s10676-010-9234-6>
- Spiggle, S. (1994). Analysis and interpretation of qualitative data in consumer research. *Journal of Consumer Research*, 21(3), 491–503. <https://doi.org/10.1086/209413>
- Streeter, T. (2011). *The net effect: Romanticism, capitalism, and the internet*. New York University Press.
- Tapscott, D. (1998). *Growing up digital: The rise of the net generation*. McGraw-Hill.
- Taylor, F. W. (1911). *The principles of scientific management*. Harper & Brothers.
- Thompson, C. J. (1997). Interpreting consumers: A hermeneutical framework for deriving marketing insights from the texts of consumers' consumption stories. *Journal of Marketing Research*, 34(4), 438–455. <https://doi.org/10.1177/002224379703400403>
- Thompson, C. J. (2004). Marketplace mythology and discourses of power. *Journal of Consumer Research*, 31(1), 162–180. <https://doi.org/10.1086/383432>
- Thompson, C. J., Locander, W. B., & Pollio, H. R. (1989). Putting consumer experience back into consumer research: The philosophy and method of existential-phenomenology. *Journal of Consumer Research*, 16(2), 133–146. <https://doi.org/10.1086/209203>
- Thompson, J. B. (1990). *Ideology and modern culture: Critical social theory in the era of mass communication*. Stanford University Press.
- Turkle, S. (2011). *Alone together: Why we expect more from technology and less from each other*. Basic Books.
- Turner, F. (2006). *From counterculture to cyberculture: Stewart Brand, the Whole Earth Network, and the rise of digital utopianism*. University of Chicago Press.
- Verbeek, P.-P. (2005). *What things do: Philosophical reflections on technology, agency, and design*. Pennsylvania State University Press.
- Verbeek, P.-P. (2011). *Moralizing technology: Understanding and designing the morality of things*. University of Chicago Press.

Verian Group. (2025). Artificial intelligence: European opinion divided. <https://www.veriangroup.com/news-and-insights/artificial-intelligence-and-the-future-of-work>

Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the frontline. *Journal of Service Management*, 29(5), 907–931. <https://doi.org/10.1108/JOSM-04-2018-0119>

Žižek, S. (1989). *The sublime object of ideology*. Verso.

Appendix A: Interview Guide

This appendix reproduces the semi-structured interview guide used to conduct the twenty-four in-depth interviews on which the empirical analysis of this thesis is based. The guide was developed in advance of data collection and was used as a structuring resource rather than a script. In the conduct of each interview, the researcher exercised the latitude characteristic of semi-structured interviewing to follow emotionally or theoretically significant threads as they emerged, to probe contradictions and ambiguities, and to invite respondents to develop their own framings, while ensuring that all five thematic sections were addressed in every interview.

The guide is organised into five thematic sections of progressive depth, opening with rapport-building warm-up questions and proceeding through open associations, personal expectations, deeper reflections, and informational ecology. Approximate timings are indicative; actual interview duration ranged from approximately thirty to ninety minutes, with most interviews lasting between forty-five and seventy-five minutes.

A.1 Pre-Interview Procedure

At the start of each interview, the researcher introduced the project as a master's thesis concerned with public perceptions of humanoid robots and artificial intelligence. The respondent was informed that the conversation would be audio-recorded for transcription purposes, and was invited to provide explicit verbal consent on the recording. Respondents were further informed that their participation was voluntary, that they could decline to answer any question, and that, in the case of non-expert respondents, they would not be identified by name in the published thesis. The deliberate choice not to brief respondents on the specific theoretical framework guiding the analysis was made to preserve the spontaneity of their articulations. No projective stimuli, video material, or written prompts were introduced at any point.

A.2 Section 1: Warm-Up and Introduction (5–7 minutes)

Goal: Establish rapport and surface spontaneous initial associations.

- Could you introduce yourself and tell me a little about your everyday life, work, or studies?
- When did you last come across something related to humanoid robots, for example, a news story, video, social media post, or conversation about robots

such as Tesla Optimus, Figure, or Boston Dynamics Atlas? What went through your mind?

A.3 Section 2: Open Associations and Feelings (10–12 minutes)

Goal: Surface raw sentiments and narratives without analytical framing.

- When the words “humanoid robot” come up — a machine that looks and moves somewhat like a human — what pictures, feelings, words, or thoughts come to mind first?
- How do you personally feel about the idea of these kinds of robots becoming more common in the world? (Probes if needed: curious, uneasy, excited, neutral, something else.)
- What comes to mind when you think about what these robots could mean for people’s everyday lives, work, or relationships, in positive ways, negative ways, or both?

A.4 Section 3: Personal Expectations and Sentiments (12–15 minutes)

Goal: Deepen emotional responses and elicit personal expectations.

- If you imagine these robots starting to appear more in everyday life over the coming years, where do you picture them appearing first, workplaces, homes, public places, care settings? How does that thought make you feel?
- How would it sit with you personally to have one of these robots around you, for example, helping at work, assisting family members, or in other situations?
- Different people have very different takes on this, some see humanoid robots as helpful tools, others as something more disruptive or even concerning. What is your own take, and what examples or reasons shape it?
- Thinking about yourself or people in similar situations to you, what personal upsides or downsides do you anticipate if these robots became widespread?

A.4.1 Optional Tailored Probes

The following probes were used selectively where the respondent’s background made them analytically productive.

- For blue-collar or service workers: In your kind of work, factory, logistics, retail, or care, how does the idea of a robot being there feel?
- For white-collar professionals or managers: In an office or knowledge-work setting, what feelings come up about robots being part of the team?
- For academics: From your area of expertise, what broader questions or feelings do these robots raise for you?
- For industry leaders or experts: From what you have seen first-hand, what stands out emotionally or practically?

A.5 Section 4: Deeper Reflections and Mixed Feelings (8–10 minutes)

Goal: Allow contradictions, evolving thoughts, and reactions to recurrent public propositions to emerge organically.

- As we have been talking, have any mixed feelings, changes in perspective, or new emotions about humanoid robots come up for you? Could you share a bit about that?
- Elon Musk has predicted that around 2040 there could be approximately ten billion humanoid robots on Earth, outnumbering the human population. What do you think about such a statement? How do you envision a future of that kind?
- Some commentators argue that work may become optional in the future, with people supported by Universal Basic Income while robots and artificial intelligence perform most labour. How would you feel about such a future?
- What kinds of experiences, stories, or sources have shaped how you feel about this topic?

A.6 Section 5: Information Sources and Closing (7–10 minutes)

Goal: Probe media and informational ecology, and provide an open closing space.

- Do you ever come across discussions, videos, news, or posts about humanoid robots or advanced robotics, for example on YouTube, TikTok, Reddit, LinkedIn, news apps, or in conversations with friends, family, or colleagues?

- If so, which sources do you tend to engage with most, and how do they affect your feelings or thoughts on the topic? (If not: That is interesting, where else, if anywhere, do you get information when something like this catches your attention?)
- Is there anything else about humanoid robots, their possible role in society, or how they make you feel that we have not touched on yet, which you feel is worth mentioning?
- Finally, would you describe yourself as overall optimistic or pessimistic about the future?

A.7 Closing Procedure

At the close of each interview, the researcher thanked the respondent for their participation, repeated the assurance of anonymity for non-expert respondents, indicated that audio recordings would be used solely for transcription purposes, and offered the option of receiving the completed thesis on request.

Appendix B: Use of Generative Artificial Intelligence (AI)

In accordance with WU Vienna's requirements for the use of generative artificial intelligence (AI) in master's theses, the following declaration is made. Generative AI tools (specifically Anthropic Claude, a large language model) were used as a writing aid in the preparation of this thesis. The use of AI was directed and supervised by the author throughout.

The generative AI tool was employed to support the following activities: (1) drafting and revising prose for clarity, concision, and stylistic consistency, including the reduction of long sentences, the removal of stylistic patterns such as em-dash overuse, and the strengthening of paragraph transitions; (2) generating preliminary drafts of connective passages between thematic sections, which were then substantively revised by the author; (3) verifying bibliographic information for cited sources against publicly available databases; (4) drafting the SVG vector graphics for Figures 1 and 2 based on specifications supplied by the author.

All theoretical claims, interpretive arguments, empirical analyses, and substantive contributions in this thesis are the author's own. The twenty-four semi-structured interviews that form the empirical basis of the analysis were conducted, transcribed, and analysed by the author. The five-cycle hermeneutic part-to-whole interpretive procedure described in Chapter 3 was carried out by the author. The cross-cutting patterns identified in Chapter 4 and Appendix F emerged from the author's sustained engagement with the transcripts. The model extensions proposed in Chapter 5 and the theoretical coda offered in section 6.6 are the author's contributions to the literature.

All AI-generated content was critically reviewed and revised where necessary, and integrated into the thesis with full responsibility taken by the author for accuracy, originality, and academic integrity. The author retains complete responsibility for the content of this thesis.

Appendix C: List of Aids Used

The following aids were used in the preparation of this thesis:

- Microsoft Word for document composition and formatting.
- Otter.ai and manual transcription for converting interview audio recordings to text.
- Anthropic Claude Opus 4.7 (large language model) as a writing aid, as described in Appendix B.
- Scholarly databases including EBSCO, Google Scholar, JSTOR, and the WU Library catalogue for literature search and source verification.
- iPhone 15 pro as audio recording device for capturing interview conversations with informed verbal consent from respondents.

Appendix E: Respondent Profiles

This appendix provides a profile for each of the twenty-five individuals who participated in the twenty-four semi-structured interviews that form the empirical basis of this thesis. The four expert respondents consented to be identified by name and are presented here under their real names. The twenty-one consumer respondents, including the two participants in the joint Hong Kong interview, are presented under the first-name pseudonyms used throughout the thesis. Profiles include the professional role, the relevant background context, and where appropriate the national or geographic location. They are ordered alphabetically by the name as it appears in the thesis.

E.1 Expert Respondents



Alexander Ietan. Chief Information Officer of an Austrian industrial construction group actively integrating robotic systems into its operations. Ietan operates at the practical interface of large-scale robotic deployment and corporate strategy, and his interview combines first-hand experience of industrial automation with a wider reflection on the place of humanoid robotics in domestic and societal life.



Emily Kate Genatowski. Doctoral researcher at the University of Vienna and former Google employee, conducting a privately funded research project on the regulatory gaps surrounding humanoid robotics. Genatowski lives and works alongside a Unitree C1 EDU2 research prototype humanoid robot named Tova. Her research focuses on data privacy and security, insurance and liability, transport regulations, identification, and labour displacement and taxation. She is the only respondent in the corpus whose engagement with humanoid robotics is direct and continuous rather than anticipated.



John Stringer. Founder of Tesla Owners Silicon Valley, one of the longest-established and most active Tesla owner communities in the world. Through more than a decade at the centre of the Silicon Valley Tesla community, Stringer has developed sustained personal relationships with senior figures at Tesla and SpaceX, including Elon Musk, whom he has interviewed on multiple occasions. This positions him uniquely close to the ideological discourse he articulates: he is not commenting on Tesla and Optimus from the outside but speaking from within the community in which the Techtopian articulation of the technology is most fully developed. His interview articulates the substantive Techtopian position in the corpus with the greatest internal consistency.



Martin Reisenbichler. Marketing and artificial intelligence specialist who has worked with low-code technologies, machine learning, and generative artificial intelligence since 2019. Reisenbichler began applied work with GPT-2 in 2019 and has since identified, tested, and deployed practical generative-AI use cases across professional contexts. His interview combines marketing-domain expertise with sustained technical familiarity with the underlying tools.

E.2 Consumer Respondents

Adrien. Respondent in his late twenties with a focus on geopolitical surveillance dynamics. Adrien's interview articulates the most developed version of the geopolitical-origin concern, with specific reference to Chinese state surveillance infrastructure and the social credit system.

Anya. Marketing manager in her early thirties living and working in London. Anya's interview articulates the city-life and professional-development implications of widespread humanoid adoption with particular clarity, and provides one of the most analytically revealing accounts of the work-purpose distinction in the corpus.

Bruno. Bachelor student at WU Vienna, originally enrolled in business administration and now focused on the technical side, including data and artificial intelligence. Bruno reads independently outside his coursework and articulates the mastery coping strategy in the most developed form in the corpus.

Charlotte. Corporate lawyer in her late twenties who completed her studies approximately two years before the interview. Charlotte is Austrian and works in the typical corporate-law environment of an international firm. Her interview foregrounds privacy, surveillance, and the regulatory implications of humanoid robotics.

Elena. Professional in finance, controlling, and reporting, with a background in business informatics. Elena's interview articulates the qualified-optimism architecture in one of its clearest forms, and her closing reflection captures the operative logic of contemporary optimism-discourse explicitly.

Felix. Marketing master's student approaching graduation at WU Vienna, with strong personal interest in technology, gaming, and media. Felix's interview is notable for its frequent reference to specific recent developments in humanoid robotics and its articulation of the velocity-anxiety pattern.

Henrik. Respondent whose articulation of negative affect about humanoid robotics is among the strongest in the corpus. Henrik's interview develops the work-displacement and meaning-collapse concerns with particular force, and his closing optimism reads as the most clearly performative in the corpus.

Ho. Hong Kong respondent in his twenties working in computer forensics and financial-crime investigation. Ho participated in a joint interview together with Maya. His professional background informs a distinctive technical perspective on the surveillance and security dimensions of humanoid robotics.

Isabel. Mother and molecular biologist by training, now working in the catalyser industry. Isabel completed her master's at Lund University in Sweden. Her interview articulates the most explicit prohibition in the corpus on robotic care for children, and her position on emotional attachment to humanoid robots is among the most resistant.

Karin. Public health and psychology researcher who completed her studies the year before the interview and worked for approximately ten months in the field at the time of the conversation. Karin explicitly disclaims technical expertise in AI or robotics, which positions her interview as one of the more representatively non-specialist accounts in the corpus.

Liliana. Big Four professional in her late twenties. Liliana's interview foregrounds discomfort with the humanoid form itself and articulates the pre-emptive-distancing coping strategy in one of its clearest forms.

Lorenz. Self-employed nutritionist and personal trainer in his thirties, with formal training in nutritional sciences. Lorenz works primarily in B2B and B2C nutritional counselling. His interview develops the historical-philosophical register on technological dependence further than any other respondent in the corpus, with explicit engagement with the arguments of Theodore Kaczynski.

Lukas. Computer scientist with a master's degree in artificial intelligence, currently working as a programmer at a fintech firm generating financial reports. Lukas combines deep technical familiarity with the underlying systems with a cautious position on domestic deployment, and his interview articulates the surveillance concern with particular technical specificity.

Markus. Chief Executive Officer of the Swiss office of an international insurance broker. Markus is forty-five years old, lives in Switzerland, and has approximately twenty-five years of experience in the insurance industry. He is personally interested in digitalisation and in Tesla, and his interview articulates the most developed version of the pro-attachment position on human-humanoid relations in the corpus.

Matteo. Marketing professional with strong personal interest in new technologies, particularly AI-generated images and video. Matteo's interview articulates the instrumentalist 'capable tool' framing of humanoid robotics with particular candour.

Maya. Hong Kong respondent in her twenties working in capital markets after a previous role at Bloomberg, with a business degree from the Hong Kong University of Science and Technology. Maya participated in a joint interview together with Ho. Her

interview provides one of the corpus's most ambivalent middle-position accounts of emotional attachment to humanoid technology.

Nadia. Marketing campaign manager in her mid-twenties working for an international consumer-goods company. Nadia is Bulgarian and Polish, grew up in Eastern Europe, is married, and has one child with another expected. Her closing optimism is the most affectively candid in the corpus.

Petar. Final-year music student and professional musician, with active engagement across performance, teaching, event management, and recording. Petar has lived in Vienna for approximately five and a half years. His interview articulates the most overtly hostile position in the corpus, with the hostility tied specifically to surveillance concerns.

Sofia. Primary school teacher at an international school, teaching seven-to-eight-year-olds. Sofia's interview articulates the qualified-optimism architecture with the greatest clarity in the corpus, and her closing reflection captures the operative logic of contemporary optimism-discourse explicitly.

Stefan. Professional at an EU delegation in Vienna to an international organisation, working in public-sector diplomacy. Stefan represents the EU's interests and those of its member states at the various UN organisations based in Vienna. His interview articulates the geopolitical-origin pattern with the most systematic country-level analysis in the corpus.

Tomáš. Master's student in marketing, part-time marketing account manager at a Viennese digital marketing agency, and part-time entrepreneur in 3D printing. Tomáš's interview is among the most theologically and philosophically resistant in the corpus, with explicit reference to religious frameworks for evaluating human-robot relations.

Appendix F: Additional Narrative Patterns

The analysis presented in Chapter 4 surfaced twelve cross-cutting patterns in the corpus. Six of these directly motivate the structural extensions to Kozinets' (2008) model developed in Chapter 5 and are therefore documented in the main body. The remaining six are recorded here. They are analytically substantive and recurrent across the corpus, but they bear on the proposed model extensions more indirectly, either by elaborating cleavages that the body patterns already establish or by operating at a meta-level relative to the ideological field itself. They are retained in full, rather than

omitted, in the interest of analytical completeness and procedural transparency. A reader assessing the adequacy of the revised model should be able to inspect the patterns that informed the analysis without becoming load-bearing for the proposed extensions. Each section below preserves the interpretive claim and the model relation recorded in the original analysis. Cross-references in Chapters 4 to 6 that point to these patterns direct the reader to the corresponding section here.

F.1 The Domestic Sphere as Ideological Boundary

The home is the sharpest cleavage in the corpus. No other contextual variable (workplace, public space, military application, or healthcare setting) divides respondents as cleanly along ideological lines as the question of whether a humanoid robot belongs in domestic spaces. Respondents who endorse robotic deployment in factories, warehouses, hospitals, and even battlefields will, with surprising frequency, refuse it for their own home. Respondents who welcome it at home tend to be the same respondents who articulate the most fully consummated Techtopian or Work Machine ideologies elsewhere in the interview. The home operates as the main threshold at which respondents' otherwise loose ideological commitments are forced into hard form.

The case against domestic adoption is articulated consistently across very different respondents. Sofia speaks for many: "At home it's so personal, and I like to keep the human experience more present there. I'd accept a robot in maybe a workspace or other areas where it enhances your life. But at home, unless there was a very specific reason, I don't think I'd want one." Tomáš draws a more analytically explicit boundary, distinguishing between robots that interact with the environment (acceptable) and robots that interact with him (not acceptable): "It would really be a robot, not a companion. It could be a gardener, a cook, a cleaner. But what I'm trying to say is that robot would interact with the environment — not with me." Isabel articulates an absolute prohibition tied to dependants: "Would I ever leave my child with a robot — absolutely not. And the same goes for an elderly person. If they need help, it's because they need human interaction." Petar, whose conditional acceptance is among the most reluctant in the dataset, allows the robot only on terms of strict invisibility: "If it's just a slave, costs ten grand, I never need to cook or clean again, and somehow somebody guarantees nothing will ever go wrong... As long as it can stand in a corner where I can't really see it... The idea of it sitting there like it's my friend? That's sick."

The case for domestic adoption is articulated with comparable conviction by a smaller and more identifiable subset of respondents, those whose dominant ideological register

is robustly Techtopian, Techspressive, or Work Machine. Alexander Ietan opens the case with the canonical formulation: “Who doesn’t want an assistant who does the tedious household tasks?” Matteo extends it to the limit, with a phrasing that itself becomes interpretively significant: “It’s like having a very capable tool. A slave, almost, if I’m being blunt. That would be convenient.” Nadia volunteers the strongest endorsement of all: “I’d trust a robot to look after the kids,” and frames her preference for an AI doctor over a human in identical terms: the robot is more reliable, less biased, and emotionally even-tempered. Markus ties the case for domestic adoption to the eldercare frame. He describes his eighty-seven-year-old father, who lives alone, “doesn’t want outside help, but has no one to speak to because all his friends and colleagues have passed away,” and identifies this population as the natural first beneficiary of a humanoid presence in the home. Emily Kate Genatowski, who already lives with a humanoid prototype, develops the most sophisticated pro-domestic argument in the corpus. She advocates not for the unmodified humanoid in an unmodified home, but for what she terms “industrialising our domestic spaces”, that is, reconfiguring the home around the robot’s capacities, much as nineteenth- and twentieth-century homes were reconfigured around electricity, gas, and the appliance.

The home is the site at which the supreme goods of Pleasure (Techspressive) and Productivity (Work Machine) are most directly contested by an inarticulate but powerful counter-good that the respondents struggle to name but consistently invoke: the home as the last reserved space for the human. The pattern affirms Kozinets’ Green Luddite element, the supreme good of preserved human ways of life, while extending its point of operation. Where Kozinets’ original respondents articulated Green Luddite positions primarily in relation to the natural environment and traditional crafts, this corpus articulates them in relation to a domesticity that is itself increasingly understood as a fragile inheritance to be defended.

F.2 The Uncanny Valley as Ideological Trigger

This pattern concerns the human form itself. A consistent subset of respondents articulate a marked and analytically interesting disjunction: they accept robotic automation in principle but react with visible discomfort to robots that adopt the human form specifically. The discomfort is not a mere aesthetic preference. It is a reliable trigger of moral and emotional complexity that non-humanoid robots do not provoke. The same respondent who is comfortable with a Roomba, a Boston Dynamics quadruped, or a sushi-bringing service robot will pull back sharply when the form factor shifts to something that resembles ourselves.

Elena names the disjunction precisely: “Not one that looks like a human — that would be scary for most of us... I think the form should be machine-like rather than human-like. Something that has the ability to do what a human does, but doesn’t look like one.” The defining feature is not capability but resemblance. Liliana takes a related position from a slightly different angle, articulating discomfort with the very terminology before any specific deployment is considered: “I don’t really like that term, to be honest. It scares me a little... It’s the ‘humanoid’ part — it’s associating something so closely with us. It’s really strange to imagine something that seems human, that maybe has all the characteristics of a human being. My big fear is the idea that in the future you might not be able to distinguish between a human and a robot.” Martin Reisenbichler, who is one of the most technologically literate respondents in the corpus, gives the phenomenon its clinical name: “The uncanny valley effect is real — if a robot looks almost but not quite human, it can be deeply unsettling.”

The most analytically revealing case is Anya’s. She describes encountering a meme on Instagram in which a humanoid robot was deliberately made to malfunction, and reports an empathetic reaction that surprises her: “I empathised a lot with the robot. When I saw it fall and stuff, I felt bad for it — the way people were laughing at it. And in that moment I thought, why am I feeling bad? At the end of the day, it doesn’t have emotions.” When asked to locate the source of the empathy, she replied: “Rationally, I know it’s not a human being. But yeah, I think with the robot it was more about the shape. It really looked very similar to us.” These conversations demonstrate the uncanny valley operating in both directions, generating discomfort in some respondents and empathy in others. In either case, the form activates moral and emotional responses that the same respondents do not extend to identically capable but differently shaped machines.

A further variant of the pattern emerges in Markus’s legal speculation. He anticipates that the human form will eventually compel a new legal category, “neither human nor thing”, analogous to the reform that, in Switzerland, transformed animals from legal property into living creatures with their own standing. The human form is therefore not just a trigger of psychological response. It could lead to institutional response, force legal, ethical, and social systems to develop categories that did not previously exist.

In Kozinets’ (2008) terms, this pattern operates as a transversal trigger. Meaning, a contextual feature that intensifies whichever ideological register a respondent already inhabits. For Tectopian respondents the human form generates anticipation. For

Green Luddite respondents it generates revulsion. For Work Machine respondents it generates the question of whether the human form is functionally optimal at all (Tomáš, in particular, asks why the human form is the most efficient solution for navigating physical environments). As a result, The uncanny valley acts less like an ideological position, but rather as an amplifier of positions already held. It is also one that the original 2008 model does not anticipate, because the technologies it studied did not pose the question of human resemblance with any urgency.

F.3 The Work-Purpose Distinction and the UBI Cleavage

This pattern concerns the relationship between work, meaning, and the proposed policy responses to large-scale automation. Across the corpus, respondents consistently distinguish between two different consequences of mass labour displacement: the loss of income, which most respondents treat as a problem in principle solvable by policy, and the loss of purpose, which a majority treat as a problem that policy cannot solve and that the available proposals do not address. The distinction is structural and recurring. The policy debate that follows from it, the debate about Universal Basic Income, is the corpus's most articulated political cleavage.

Alexander Ietan poses the underlying question in its most direct form, immediately after observing that a great deal of his own professional domain is automatable: "What is human purpose if you're not contributing through work? That's the deep question we need to ask, and I'm not sure I have a good answer." Markus develops the position systematically: "Doing nothing is dangerous. People will become bored and develop ideologies, even abuse drugs or alcohol... Work is identity, work brings normality, work brings discipline. We are made for it." Matteo articulates the same anxiety in personal terms: "What scares me most is not the lack of money — it's the lack of meaning. If I cannot use my time meaningfully and feel productive, my well-being deteriorates fast." Nadia extends it to societal scale: "If you remove work from people, especially in cultures and generations that grew up with the idea that work gives you meaning, then they're going to be miserable." Lorenz adds the historical-philosophical register, citing Theodore Kaczynski and explicitly disclaiming his methods while endorsing the underlying concern. The articulations converge on a single substantive claim: meaning is constitutively tied to productive contribution, and its removal is therefore not a neutral expansion of leisure but a normatively significant loss. This appears to be the work-purpose equation the proposed UBI policy framework, in the eyes of these respondents, fundamentally fails to recognise.

A minority position pushes against this consensus, sometimes sharply. Matteo simply does not share the premise that work is a precondition of meaning: "Probably full of people enjoying life. Having fun. I think most people don't like their jobs — I do, I enjoy mine — but I think most people would prefer to stay home." Felix offers the most theoretically optimistic version, invoking the classical world: "I think about the golden age of Greek philosophy, where certain people could spend their whole day just thinking about the world. Something like that could be very productive for humanity." Nadia takes a structurally similar but more pragmatic position: "I don't fear it — I'd just find something else. The boring jobs are the ones getting replaced anyway, and I'd be happy to move towards something more creative and intellectually stimulating." Emily Kate Genatowski is openly pro-UBI, but couples her endorsement to a substantive critique of current taxation arrangements that she argues distort the labour market in favour of automation: "We need to start taxing the value added by labour rather than the labourer themselves. If a robot is replacing a worker, there should be an economic incentive for the business to either not do that or to pay into a system that helps that person reskill or transition."

The structure of the cleavage is informative in itself. The respondents who oppose UBI tend to oppose it not on cost or feasibility grounds (those concerns are mentioned but rarely lead the argument) but on anthropological grounds: a claim about what human beings are and what they need. The respondents who endorse UBI tend to endorse it not because they disagree with that anthropology but because they identify an alternative source of meaning (creative pursuits, philosophical reflection, leisure-as-flourishing) that they believe will substitute for the meaning currently supplied by work. Both positions accept the work-purpose distinction the pattern names; they disagree about whether purpose can be reconstituted under conditions of mass non-employment, and about what role policy should play in attempting it. Reskilling and retraining emerge across the corpus as the more frequently endorsed intermediate solution, supported by Lukas, Stefan, Charlotte, Liliana, Tomáš, and several others. They tend to be preferred precisely because they preserve the structural relationship between effort and reward that respondents associate with meaning-bearing activity.

In Kozinets' (2008) terms, this pattern develops the inner contradiction of the Work Machine element with unusual specificity. Kozinets identifies the Work Machine articulation of technology with the supreme good of Productivity and notes that the element's shadow side is the loss of meaning, the dehumanisation of work, and the alienation that critical theorists from Ellul to Heidegger have analysed. The present corpus shows that those shadows are now articulated at the centre rather than the

periphery of the Work Machine ideology. Respondents do not merely register them as theoretical concerns; they cite them as the operative reasons for their resistance to specific policy proposals, and they develop them through reference to concrete personal or generational experience. The pattern affirms the model's theoretical scaffolding while moving the inner contradiction from the margins to the centre of the analytical attention the element receives.

A further analytical axis runs through the work-purpose articulations gathered above that the present analysis was not originally calibrated to surface but that, on examination, the data support. The respondents who articulate the loss-of-purpose concern in its most existentially loaded form, that is, those who frame the question of automation primarily as a threat to what work is for rather than as a problem of income or capability, are predominantly male. Alexander Ietan's foundational formulation ("What's the meaning of life if work disappears?"), Tomáš's structural argument about purpose and meaning, Stefan's reading of Musk and Bezos as evidence that meaning exceeds money, Bruno's anxiety about non-pecuniary fulfilment, and Henrik's extended reflection on the unemployed friend all locate the analytical centre at the meaning function of work. Of these articulations, only Charlotte's normative warning comes from a female respondent in the corpus, and her articulation, while substantively aligned, is the most pragmatic of the group. The female respondents in the corpus tend to articulate the work-and-meaning question with greater equanimity. Nadia would "just find something else." Anya's closing posture toward generational adaptation does not pass through the meaning-of-work register. Sofia's anchoring concern is relational rather than vocational. The structurally similar position to Nadia's is also occupied by Matteo, who is male, so the pattern is gendered but not gender-determined.

The pattern is consistent with a substantial sociological literature on gender, work, and identity. Connell's (1995) account of hegemonic masculinity identifies the breadwinner role as one of the most stable and culturally enforced anchors of male identity in industrial and post-industrial societies, and shows that disruption to that role produces distinctive affective responses (anger, withdrawal, projection of blame onto outside agents) that map closely onto the more pessimistic articulations in the present corpus. Burbano, Padilla, and Meier (2024) document that male respondents place significantly more weight than female respondents on the meaning derived from work, and that this differential intensifies under perceived occupational instability. Jahoda's (1982) classic account of unemployment as social-psychological deprivation extends the argument further: the loss of work deprives the individual not only of income, but of

structured time, collective purpose, social validation, and routinised activity (see also Gonalons-Pons & Gangl, 2021; Munsch, Rogers, & Yorks, 2016).

The implication for the present corpus is that the loss-of-purpose articulation may be ideologically over-represented by male respondents not because the question of meaning is gendered as a question, but because the anticipated subjective consequence of work disappearing is heavier for respondents who occupy traditionally male-coded labour-and-identity arrangements. The framing is not deterministic and admits of significant variation within the corpus. Charlotte's articulation is closest to the male cluster's, and Nadia's distance from it is in part a function of her professional position rather than her gender alone. The gradient is nonetheless consistent enough to warrant explicit registration. The work-purpose distinction may, on the reading the corpus supports, be a distinction whose ideological force is partly gendered in its present cultural moment. The policy debate about Universal Basic Income may accordingly land differently with male and female respondents not because they disagree about what work is for, but because they anticipate different subjective consequences of its disappearance.

This observation has implications for the model extension developed in Chapter 5. The work-and-meaning cleavage may be a particularly clear instance of a more general phenomenon: the inner contradictions Kozinets (2008) identifies within each element are not equally distributed across the population, and the structural location of the respondent, including their gender, generation, profession, and family role, partly conditions which contradictions surface as analytically operative and which remain latent. The present study cannot pursue this observation systematically because the sample is not stratified to support gender-comparative inference at the level of statistical reliability. The pattern's robustness across the male respondents in the corpus is nonetheless sufficient to register the gender axis as an analytical dimension that future quantitative work could productively pursue.

F.4 The Emotional Attachment Split

This pattern names a cleavage that runs through the corpus on the question of human-robot emotional bonds. Respondents do not disagree about whether such bonds will form: there is an unusual degree of consensus that they will. They disagree, sharply and along recognisable ideological lines, about what to make of that. The cleavage divides respondents into those who frame human-robot attachment as a natural and largely benign extension of existing human capacities, and those who frame it as a

symptom of a deeper social pathology that the technology will worsen rather than alleviate.

The pro-attachment articulation is most fully developed by Markus, whose argument operates on three levels at once. At the personal level, he invokes his eighty-seven-year-old father, alone, friendless after the deaths of his peers, refusing outside help, as the natural beneficiary of robotic companionship. At the cultural level, he extends the argument through the hairdresser analogy: "Everyone in town goes to the hairdresser and speaks for an hour about what's going on in their life. Imagine if a robot could play that role — and at the same time detect early signs that someone might be heading towards depression or a mental health issue, and either help them directly or connect them with the right professional." At the empirical level, he cites the existing pattern of human-bot bonding in gaming environments: "people form genuine emotional bonds with bots in World of Warcraft and similar environments. It's human nature." The composite argument frames human-robot attachment as continuous with longstanding human capacities for parasocial connection, and as a potentially significant intervention in the contemporary mental-health crisis. Matteo, asked whether people will form bonds with robots, treats the question as practically settled: "People are already getting attached to chatbots. Of course they're going to form bonds with a physical humanoid robot." Alexander Ietan develops a similar position with reference to the trust that builds rapidly when the technology works: "A lot of people will trust them quickly, give them many tasks, and maybe even see them as friends."

The contra-attachment articulation is developed with comparable conviction by a different cluster of respondents. Tomáš offers the most analytically sustained version. He distinguishes sharply between the kind of relationship in which bonding takes place and the kind of relationship that bonding requires: "A partner is someone who challenges your thoughts, has a different opinion, a different worldview. The AI girlfriend products that are popular now are not models that challenge you — they're models that make you feel good. That's not beneficial in the long term, even for the individual." Charlotte articulates the same concern in terms of social-media analogy: "I think it's comparable to people becoming very invested in social media — it's also not a real connection with another human. We need to be careful that people don't end up living entirely in their own bubble." She identifies what she takes to be the operative mechanism. AI tends to agree with the user, the user's closest human relationships will continue to push back, and if the AI relationship displaces the human ones, the user is deprived of honest feedback. Elena arrives at the same diagnosis

from a different starting point: "It's super dangerous to believe it's a real person. You have to know your limits — what's real, and what to expect from a tool. It cannot replace feelings or a human." Sofia concludes with the same proscription: "They're not humans. Even if we want them to seem more human, they're not."

Lorenz articulates perhaps the most historically resonant version of the concern through an unusual analogy. AI-mediated emotional interaction, he argues, is operating in the same structural register as pornography: "an easier, lower-friction version of something that exists in reality, and people gravitate towards the easier path." The analogy frames the AI partner not as a substitute for the absent human partner but as a competitor to the present one, and one whose competitive advantage is precisely the absence of the friction that makes real human relationship work. Maya, offering a more ambivalent middle position, recounts a friend in Canada who relies on ChatGPT for emotional support after migration. The use is comprehensible and locally functional, but her broader concern is that prolonged reliance "will have a prolonged effect on your interpersonal skills, your confidence around people, how you grow as a person." Even Anya, who is otherwise relaxed about most aspects of the technology, concedes the point. She has friends who joke about marrying ChatGPT, and although she treats the joke as a joke, she recognises the underlying social dynamic: "unfortunately there will be people who feel very alone and end up closing themselves off even more."

The cleavage maps with reasonable precision onto Kozinets' (2008) structural division between collectively oriented and individualistically oriented ideological positions. The pro-attachment respondents tend to articulate their case in terms that prioritise individual welfare and the alleviation of specific personal sufferings, that is, the lonely father, the isolated migrant, the user under emotional stress. They treat the broader social consequences as either downstream concerns or as outside the scope of the immediate moral question. The contra-attachment respondents tend to articulate their case in terms that prioritise the integrity of the social fabric: relationship as honest mutual challenge, social-media-style filter bubbles, the cumulative effect on interpersonal capacity at the population level. The cleavage is not principally about the technology. It is about whether the appropriate ethical unit of analysis is the suffering individual or the social formation. In Kozinets' (2008) terms, this is a productive instance of the contrariety of individualism that the model identifies. The unusual feature is that both sides of the cleavage are developing variants of what the model would treat as Techtopian articulations, differing not on the supreme good but on its appropriate scale.

F.5 The Generational Gradient

This pattern names a structural feature of the corpus that organises a great deal of the variance observed in the other patterns. Respondents repeatedly invoke, anticipate, or describe a generational gradient in the reception of humanoid robotics: the expectation that children will normalise what adults find troubling, and that the moral and emotional baseline against which the technology will eventually be assessed is one that has not yet been established. The gradient is not merely a demographic observation. It is a recurring rhetorical resource through which respondents frame both their hopes and their concerns about the technology.

Emily Kate Genatowski offers the cleanest empirical observation, from her year of public encounters with her own humanoid prototype: “Older people come up and speak to me about it; younger people come up and speak to it directly.” The grammatical distinction is interpretively heavy. To speak about the robot is to treat it as an object of conversation. To speak to the robot is to treat it as a participant in conversation. The generational divide, in her account, is not principally about attitude or opinion but about the basic categorisation of the entity in question.

Nadia provides a domestic counterpart from her life as a parent of a young child: “In the Chinese restaurant nearby, they have a serving robot — just a stack of plates with a monitor on top showing a smiling cat. My daughter loves it. The first time she saw it she was taken aback, but she’s taken aback by new humans too. Then she gets genuinely excited about interacting with it. The human-robot barrier that we feel now will probably just dissolve for future generations.” The disappearance of the human-robot barrier for the child is figured as a prediction about what future adult populations will look like, and, implicitly, as a partial reassurance about the technology, since the resistance the present generation feels is identified as a generational artefact rather than a permanent moral position. Markus reaches the same conclusion through gaming culture, which he treats as the leading indicator of the next generation’s emotional repertoire. The bonds people already form with non-player characters in massively multiplayer online environments are, in his reading, the natural prologue to the bonds his children’s and grandchildren’s generations will form with embodied AI systems.

The generational gradient is not always invoked optimistically. Elena cites the post-millennial intelligence trajectory she has read about and identifies AI use among children as a contributing factor. Adrien warns that the dependency he sees forming in the youngest cohort will, in the event of any infrastructural disruption, leave them severely under-equipped. Isabel articulates the position from the maternal side with

maximum sharpness (“Would I ever leave my child with a robot — absolutely not”) and identifies the generational shift as a danger to be resisted rather than a destiny to be welcomed. Anya, whose life sits between the two cohorts she compares, reverses the standard direction of generational anxiety: “I have good hopes for the younger generations. I think our generation got the worst of it because it was all so new and we didn’t know how to use it properly... the younger generations seem better at using what they have in a healthier, more positive way.”

Several respondents invoke their own children in ways that read as ideological positioning more than as biographical disclosure. Alexander Ietan places his nine- and eleven-year-old children at the centre of his closing optimism: “I’m not sure they’ll even need a driving licence in eight to ten years. They’ll grow up with full exposure to robots in daily work and daily life.” He ties his hope for the technology specifically to the prospect of robotic care for himself in old age. John Stringer frames his parental responsibility as the cultivation of the dispositional capacity to handle the change: “As a father of four, I’m not going to stick my head in the ground and wish this wasn’t happening... I’m going to do my best to train my kids to be as curious as possible. That curiosity — that’s what prepares people for whatever the future brings.” Matteo, by contrast, registers the parental concern in negative form. He would not know what to advise his son to study, given the uncertainty about which careers will survive the next decade.

A particularly evocative version of the gradient is offered by Stefan. He describes his mother, initially highly sceptical of all digital technology, including online banking, taking up a philosophy degree late in life, gaining access to ChatGPT through a family member, and now using it daily in ways that have transformed her relationship to the broader technological landscape. The arc he describes (initial scepticism, exposure, positive personal experience, normalisation) is presented as the template he expects to see at scale. It is not a generational shift in the strict sense but an exposure-driven adoption arc that, given enough time, produces something resembling generational difference even among individuals.

In Kozinets’ (2008) terms, the generational gradient operates as a temporal extension of the model’s relational logic rather than as a new ideological element. Respondents do not articulate the generational dimension as an ideological position in itself. They invoke it as a way of reasoning about the trajectory of the field, about where the dominant articulations will sit in twenty years, about whether their own resistance or enthusiasm is the leading or trailing edge of the eventual settlement. This is, on the present study’s reading, a productive supplementation of the model. The field is not

only configured at any given moment but is understood by those who inhabit it to be moving along a generational vector whose direction is reasonably predictable even where its specific destination is not.

F.6 The Robot as Social Mirror

This pattern operates at a meta-level relative to the others documented in this study. Across the corpus, respondents repeatedly identify, sometimes about themselves, more often about others, the phenomenon by which the humanoid robot functions as a projective surface. The entity onto which the speaker's prior ideological, emotional, political, and even sexual content is mapped, is then read back as if it were a property of the entity itself. The pattern is reflexive and meta-analytical, and it has the unusual feature of being articulated most clearly by respondents whose direct engagement with the technology is the most extensive.

Emily Kate Genatowski offers the canonical formulation. Asked about the public reactions her humanoid prototype provokes, she responds with what is, on this study's reading, the single most analytically significant sentence in the entire corpus: "What fascinates me is that the robot acts as a prism. Whatever someone feels about the future — fear, excitement, hope, threat — comes pouring out when they stand in front of it." She extends the observation through her romantic life. Of two partners she has had since acquiring the robot, one read it as a competing partner and the other read it as a child, neither of which she had encouraged or suggested. Her summary is unambiguous: "Whatever the robot elicits in people is entirely based on what's already inside them — whether that's fear, hope, or something else. It brings out whatever is there."

The pattern is corroborated, in different registers, across the corpus. Markus anticipates the consumer choice of robots through the analogy of avatar customisation in video games: "People will choose humanoid robots the same way: the look, the gender, the personality. And we're still human — if my wife is laughing with a robot that looks like a young Brad Pitt, I'm probably not entirely comfortable with that either. Even with robots, we bring all our very human insecurities." The robot, in his reading, is not merely chosen for the projection it enables. It is the consumer category through which the projection becomes a market. Matteo's spontaneous framing of the robot as "a slave, almost" performs the same operation in less reflexive form. The framing is not an inference from the robot's properties but a projection from the speaker's prior conceptual repertoire about labour, ownership, and convenience. Petar's elaborate speculative narrative about Tesla taxis taking citizens to police stations, robot officers

replacing human ones, and government-controlled humanoid populations in the streets reads, on close interpretation, less as a prediction about the technology than as an articulation of the speaker's prior political fears for which the humanoid robot supplies the figurative material.

Adrien articulates the pattern in its most morally explicit form: "Human nature... that's what worries me. The technology itself isn't the villain. We are, potentially." The reframing is significant. The conventional moral grammar of the technology debate locates agency and responsibility in the artefact (or in its developers). Adrien shifts the locus back to the user, and the consequences of the technology are reframed as the consequences of what humans choose to project into and onto it. Tomáš develops a structurally identical move from a different starting point. The AI girlfriend phenomenon is, in his reading, less a fact about the AI than a fact about the loneliness and unmet need for which the AI is the available container. Lukas puts the cognitive science under the moral observation. Humans are evolved to attribute consciousness and intent to anything that exhibits surface signals of either, and the humanoid robot is precisely the object designed to maximise such signals. Whatever the user reads back from it is, in this account, almost necessarily an artefact of the user's own cognitive architecture rather than a property of the device.

The pattern is double-loaded analytically. On the first reading, it is a finding about consumer responses to humanoid robotics: respondents project, and the analyst can document the projection. On the second reading, it is a finding about the reflexive sophistication of the present corpus: respondents themselves recognise the projection, name it, and incorporate its recognition into their own articulations. The reflexivity is not uniform. The most heavily projected articulations in the corpus, such as Petar's, are also among the least self-aware about the projection. It is nonetheless sufficiently widespread to constitute a feature of the data rather than an inference of the analyst.

In Kozinets' (2008) terms, this pattern operates at a level the model does not directly thematise. The 2008 framework analyses the ideological field as a structured space of positions about technology; the present pattern observes that respondents, in the act of articulating ideological positions, frequently use the humanoid robot as a screen onto which they project their existing positions about humanity itself. The robot appears to be interpretively saturated before it is ideologically positioned. This points toward an extension of the model in the direction of meta-level reflexivity that Chapter 5 develops.

F.6.1 The Interview as Symptom: A Search for Authority

A reflexive feature of the corpus that operates at the meta-level of the analysis itself, rather than as a substantive pattern of the kind documented above, warrants explicit registration. Across a substantial number of conversations, respondents turned the question back on the interviewer: "What do you think?" Variants of the phrase recur ("How do you see it?", "Where do you stand on this?", "What's your sense?") and were articulated by respondents whose substantive ideological positions span the entire range of the corpus. The pattern is not reducible to ordinary conversational reciprocity. It tends to emerge at the points in the interview where the respondent has articulated significant uncertainty and is searching for a resource through which to stabilise it.

The reflexive observation that the question itself constitutes data has analytical purchase. The interview is, in a hermeneutic frame, a dialogic encounter (Thompson, 1997; McCracken, 1988; Kvale & Brinkmann, 2009). The asymmetry it institutes (interviewer asks, respondent answers) is a working convention rather than a metaphysical constant, and the moments in which that asymmetry is reversed are revealing. In this corpus, the reversal consistently follows the articulation of a concern the respondent does not feel equipped to resolve through their own interpretive resources alone: the cognitive atrophy concern after Tomáš's extended reflection, the geopolitical concern after Stefan's outline of the bloc structure, the future of work after Henrik's most extended reflection on purpose, the human-form question after Liliana's articulation of unease. The respondent has reached the limit of their own analytical confidence and reaches outward for a stabilising figure. The only stabilising figure available in the room is the interviewer.

The observation is theoretically suggestive. Bauman's (2000) analysis of liquid modernity identifies the chronic structural condition of contemporary social life as a condition in which solid orientations have dissolved without being replaced by new ones, requiring individuals to construct and bear the weight of their own interpretive frameworks under accelerated change. The respondents who conclude by asking the interviewer "but what do you think?" appear to be enacting exactly this Baumanian condition. The substantive content of their concerns is articulated; the interpretive framework within which those concerns should be evaluated is what they reach for. They reach for it not because they lack the capacity to construct it, but because the cultural moment offers no consensual framework into which the reflexive work of evaluating humanoid robotics can be deposited. This is consistent with Norris and Inglehart's (2019) analysis of cultural backlash as structurally enabled by precisely this orientational vacuum.

The pattern is not univocal across the corpus. Several respondents (John Stringer, Emily Kate Genatowski, Alexander Ietan, Bruno) articulate their positions with sustained confidence and do not seek the interviewer's view. These respondents share a structural feature. They occupy professional or community positions that have already supplied them with an interpretive frame for the technology, and the resulting confidence is itself a form of authority-grounded orientation rather than an absence of it. The respondents who ask the interviewer's view tend to be those for whom no institutional resource has yet supplied a stabilising frame. The interview itself, by virtue of being conducted by someone identifiably engaged with the topic, momentarily occupies the vacant position. The phenomenon is, in this sense, isomorphic with the broader trend the political-science literature documents. Under conditions of insufficient legitimate authority, the available proximate figure is loaded with the orientational function that the absent institutional structure cannot perform.

In Kozinets' (2008) terms, this pattern appears to operate at a meta-level relative to the ideological elements rather than within any one of them. It is a feature of how the field is engaged rather than of what positions the field contains. Kozinets does not foreground this meta-level dynamic, and the 2008 model would not predict it. The model assumes that consumers occupy ideological positions and move between them; it does not assume that consumers, during the very act of articulating those positions, look outside the field for an authoritative perspective on how the field should be navigated. The empirical observation pushes the model in the direction of a more reflexive account of consumer ideological agency. Section 5.5 returns to this observation as part of the broader question of whether the model itself needs to be supplemented by an account of the meta-level reflexive dynamics that the present corpus surfaces.