



How Will Artificial General Intelligence Reshape the Early-Stage Startup Landscape?

A Qualitative Study of Founders, Investors, and Ecosystem Actors in the San Francisco Bay Area, 2025-2026

Simon CASTELAIN

Student ID: B00805634

ESSEC Global BBA

May 2026

Table of Contents

1	Front Matter.....	6
1.1	Executive Summary	7
1.2	Acknowledgements	8
2	Chapter 1: Introduction	9
2.1	1.1 The night the ground moved	9
2.2	1.2 The empirical gap.....	10
2.3	1.3 Where the study sits in the conversation.....	11
2.4	1.4 Research question and objectives.....	11
2.5	1.5 How I studied this	12
2.6	1.6 Contribution and roadmap.....	12
3	Chapter 2: Literature Review	14
3.1	2.1 AI and AGI as general-purpose technology and external enabler	14
3.1.1	2.1.1 The GPT frame and its historical clock.....	14
3.1.2	2.1.2 Digital entrepreneurship and the external enablement turn	15
3.1.3	2.1.3 The AGI definitions problem	16
3.1.4	2.1.4 Diffusion speed and the saturation condition.....	17
3.1.5	2.1.5 Closing reflection	17
3.2	2.2 Resource-based view, moats, and the commoditisation problem	18
3.2.1	2.2.1 The VRIN framework at the limit	18
3.2.2	2.2.2 Platform ecosystems and the tenant problem.....	19
3.2.3	2.2.3 Complementary assets and dynamic capabilities	19
3.2.4	2.2.4 Competence-destroying discontinuities, compressed	20
3.2.5	2.2.5 Closing reflection	20
3.3	2.3 Entrepreneurial judgment and Knightian uncertainty	21
3.3.1	2.3.1 Knight, McMullen-Shepherd, and the uncertainty core.....	21
3.3.2	2.3.2 The Townsend programme and computability.....	21

3.3.3	2.3.3 Effectuation, discovery, and the cost-of-experimentation shift	22
3.3.4	2.3.4 Judgment, skill, and AI-augmented action.....	23
3.3.5	2.3.5 Uncertainty migration as the load-bearing claim	23
3.3.6	2.3.6 Closing reflection	23
3.4	2.4 Venture capital under uncertainty: herding, horizons, and AI hype	24
3.4.1	2.4.1 Reputational herding and why it is rational	24
3.4.2	2.4.2 AI hype, valuation, and the funding-gap data	24
3.4.3	2.4.3 Hot-market dynamics and the power-law	25
3.4.4	2.4.4 AI at the other end of the table	26
3.4.5	2.4.5 Timeline misalignment.....	26
3.4.6	2.4.6 Closing reflection	27
3.5	2.5 The democratisation-versus-concentration tension.....	27
3.5.1	2.5.1 Everyday entrepreneurship meets winner-take-most	27
3.5.2	2.5.2 The democratisation evidence	27
3.5.3	2.5.3 The concentration evidence.....	28
3.5.4	2.5.4 Institutional isomorphism and the ethics silence.....	28
3.5.5	2.5.5 The tension as framework, not either-or	29
3.5.6	2.5.6 Closing reflection	29
3.6	2.6 Gap statement and transition	29
4	Chapter 3: Methodology.....	31
4.1	3.1 Philosophical positioning	31
4.2	3.2 Research design.....	32
4.3	3.3 Sample and recruitment.....	32
4.4	3.4 Interview conduct and data quality	34
4.5	3.5 Analytical approach.....	36
4.6	3.6 Member checking (April 2026).....	36
4.7	3.7 Trustworthiness (Lincoln and Guba 1985).....	37

4.7.1	Credibility.....	38
4.7.2	Transferability	38
4.7.3	Dependability	39
4.7.4	Confirmability	39
4.8	3.8 Reflexivity and positionality	40
4.9	3.9 Ethics	41
5	Chapter 4: Findings	42
5.1	4.1 Theme 1: The Resource Pivot, from Execution to Judgment	43
5.2	4.2 Theme 2: The Founder Identity Rewrite	46
5.3	4.3 Theme 3: The Timeline-Strategy Nexus	49
5.4	4.4 Three Counterintuitive Patterns	52
5.4.1	4.4a Geographic Centralisation Intensifies	52
5.4.2	4.4b Defensive Innovation	53
5.4.3	4.4c The Ethics Silence	54
5.5	4.5 Thematic Synthesis	55
6	Chapter 5: Discussion.....	57
6.1	5.1 Answering the research question.....	58
6.2	5.2 Dialogue with the literature.....	59
6.2.1	5.2.1 External enablement at the saturation limit.....	59
6.2.2	5.2.2 VRIN, relocated	59
6.2.3	5.2.3 GPT cycles, compressed	60
6.2.4	5.2.4 Uncertainty migration confirmed	60
6.2.5	5.2.5 Herding, priced.....	61
6.2.6	5.2.6 One reflection before moving on	61
6.3	5.3 The ASI horizon and the fragmenting consensus.....	62
6.4	5.4 The Bifurcated Ecosystem Model, formalised.....	63
6.4.1	P1: Enabler saturation	63

6.4.2	P2: Archetype-timeline alignment	64
6.4.3	P3: Technical-skepticism gradient	64
6.4.4	P4: Defensive-innovation advantage.....	65
6.4.5	P5: Ethics-velocity decoupling.....	65
6.4.6	P6: Relational-capital concentration	65
6.4.7	A note on the set.....	66
6.5	5.5 Complicating sources	66
6.6	5.6 Contribution summary.....	67
7	Chapter 6: Practical Implications	68
7.1	6.1 For founders	69
7.2	6.2 For investors.....	70
7.3	6.3 For accelerators and ecosystem support institutions	72
8	Chapter 7: Conclusion.....	73
8.1	7.1 Returning to the question	73
8.2	7.2 Three contributions	74
8.3	7.3 Limitations	74
8.4	7.4 Future research	75
8.5	7.5 Closing	76
9	Bibliography.....	76
10	Appendices	85
10.1	Appendix A: Interview Guide	85
10.1.1	Block 1: Current work and AI impact.....	86
10.1.2	Block 2: AGI definition and timeline.....	87
10.1.3	Block 3: Capital, talent, and strategy response	87
10.1.4	Close.....	88
10.2	Appendix B: Consent Form Template	89
10.3	Appendix C: Participant Sampling Matrix (Table 1).....	91

10.4	Appendix D: Member-Check Responses (P1, P3, P8).....	92
10.4.1	D.1: From P1, April 2026.....	92
10.4.2	D.2: From P3, April 2026.....	93
10.4.3	D.3: From P8, April 2026.....	94
10.5	Appendix E: Reflexive Journal Excerpts	95
10.5.1	E.1: The moment the paradox framing clicked	95
10.5.2	E.2: P7 and the guardedness problem	95
10.5.3	E.3: The P1 AGI moat anecdote after the April member check	96
10.5.4	E.4: Holding the Orchestrator-Humanist typology under pressure.....	96
10.5.5	E.5: The P8 remote-experiment update.....	97
10.6	Appendix F: AI Tool Usage Log.....	97
10.7	Appendix G: Coding Structure and Analytical Pivots	98
10.7.1	G.1: Coding Evolution	98
10.7.2	G.2: Analytical Pivots Log.....	103

1 Front Matter

Pagination-dependent elements (Table of Contents, Table of Figures and Tables) are deferred until the body is paginated. Only the Executive Summary and Acknowledgements sit here.

1.1 Executive Summary

The question driving this dissertation is how Artificial General Intelligence will reshape the early-stage startup landscape across the 2025-2030 horizon on which seed-stage bets actually resolve. The empirical core is ten semi-structured interviews conducted in the San Francisco Bay Area between October 2025 and January 2026, with founders, investors, and ecosystem actors, followed by member-check conversations with three of those participants in April 2026.

Analysis followed Braun and Clarke's (2006) six-phase reflexive thematic procedure, with 127 initial codes developed in NVivo and trustworthiness framed through Lincoln and Guba's (1985) four criteria. Of the ten interviews, nine were audio-recorded, transcribed in Otter.ai, and hand-corrected; the tenth participant declined recording and was reconstructed from contemporaneous notes inside a two-hour window. Thematic saturation was already visible by the seventh conversation and was confirmed at the ninth. Triangulation pulled in Stanford HAI (2025; 2026), the Anthropic Economic Index (2025; 2026a; 2026b; 2026c), Carta (2026), IMF WP/25/76 (2025), and METR (2025).

The central finding is a paradox that I spent several weeks trying to resolve before I realised it was itself the answer: the same technology is lowering the cost of entry into entrepreneurship and, at the same time, stretching the distance between the base of the ecosystem and its top. "Never easier to start, never harder to win," as one accelerator director put it. Four themes give that paradox its structure. The first is a Resource Pivot in which technical execution has been commoditised and judgment (what founders named "taste" in the first round and "discipline" on re-interview) has become the scarce input. The second is a Founder Identity Rewrite, with a typological split between Orchestrators who direct an AI-augmented team and Humanists who insist on holding the strategic call themselves, against investors who do not agree on which of the two they should be funding. The third is a Timeline-Strategy Nexus in which the deeper a participant's technical fluency the longer their AGI horizon, while investors cluster their estimates around their fund lifecycle rather than around the behaviour of the tools. The fourth is a trio of counterintuitive patterns: geographic concentration in San Francisco has tightened rather than dispersed; experienced founders deliberately cap their AI dependency as a defensive posture; and substantive ethical reflection is almost entirely absent from founder talk.

The theoretical contribution is the **Bifurcated Ecosystem Model: A Framework for External Enabler Saturation**, an extension of Davidsson, Recker and von Briel (2020). The model is formalised as six testable propositions, each paired with a falsifiability condition, covering enabler saturation, archetype-timeline alignment, a technical-skepticism gradient, defensive-innovation advantage, ethics-velocity decoupling, and relational-capital concentration. The model is positioned for future quantitative validation via a survey instrument.

The practical implications speak to three audiences. Founders are advised that proprietary data and model-agnostic architecture are the durable moats, and that the question “what is your AGI moat?” is best read as a diagnostic of investor fit. Investors are invited to treat the timeline disconnect between technical founders and generalist partners as a priced risk factor, to rebalance explicitly between Orchestrator and Humanist bets, and to ask founders to walk them through three recent decisions in which AI was deliberately not used. Accelerators and support institutions are encouraged to move from technical mentorship to judgment and taste, to treat geographic concentration as a feature to work with rather than around, and to invest in storytelling and conviction as the new bottleneck, consistent with the Winter 2026 cohort data reported by one of the participants.

The study’s limitations are honestly named: San Francisco-only sampling, survivor bias, temporal specificity at the AGI-to-ASI fault line, three-of-ten member checks, and a retrospectively reconstructed reflexive journal.

1.2 Acknowledgements

To Franck Rouault, my dissertation advisor, whose feedback on the earlier deliverables pushed me harder than I expected to be pushed.

To the ten participants, who I cannot name here and who gave me more of their time and candour than I had any right to expect. Three of them came back in April 2026 when I asked for a second round. That second round changed the analysis. They will recognise their own voices in these pages.

To the UC Berkeley network that opened most of the doors. Thirty-five cold LinkedIn messages produced one interview. A single warm introduction often produced three.

To my family, who accepted that a final-year dissertation written from San Francisco, Paris, and planes in between would take the shape that it took. And to the friends who read drafts I was not ready to show anyone.

2 Chapter 1: Introduction

2.1 1.1 The night the ground moved

Berkeley, early 2025. I was sitting in a Haas study room with a student I had met a week earlier, watching him build a B2B SaaS prototype from scratch. He opened Cursor, pasted in a Notion brief of what he wanted, and worked through the night with an AI coding assistant as a second pair of hands. By the time I left, around two in the morning, the product was live on Vercel with a working authentication flow, a Stripe sandbox, and a landing page his roommate had generated in Figma with another model. One night. One person. A year earlier, that would have been a semester's worth of work for a team of three.

I remember walking back down Telegraph and feeling something I did not have a word for yet. A few years earlier I had run a small sneaker business called Softsneakers, from 2020 to 2023, which reached about ninety thousand euros of revenue before I wound it down. Three years of unglamorous grinding: sourcing, customs, returns, small unit economics. What I had just watched in that study room was not a different kind of work, it was a different kind of cost structure. The feeling was not admiration, exactly. It was closer to recognising that the terrain I had been learning how to walk on had quietly rearranged itself.

That feeling is the origin of this dissertation.

Since then I have kept one foot in each side of the question the Berkeley night opened. In summer 2025 I worked as a Founder's Associate on a €2.85M leveraged buyout at Mojo Technologies (Ski-Mojo), where I sat close enough to the deal to see how capital structures bend around a small industrial business. Since January 2026 I have been a VC analyst at Neo Founders and One Green in Paris, screening thirty to forty seed and Series A deals a month, writing investment-committee memos, and watching the same AI-era patterns that my interview participants described show up again, in weekly deal flow, under different names. I write this dissertation as a participant in the thing I am studying, not as an outside observer. That is a bias I will return to in the methodology. It is also the reason the question feels urgent to me and not only academic.

2.2 1.2 The empirical gap

Three numbers anchor the surface picture. US private AI investment hit \$285.9 billion in 2025, up from \$109.1 billion the year before (Stanford AI Index 2025, 2026). Carta (2026) reports that 36.3% of new US startups are now solo-founded, while those same solo founders capture only 14.7% of priced-equity capital. And The Letter Two (2026), corroborated by SF Examiner (2026) Q1 data, puts roughly 60% of global AI venture funding and 75% of US Q1 2026 startup funding in the San Francisco Bay Area. Capital is doubling, founders are multiplying, the geography is concentrating, all at once.

What the numbers do not tell us is how the people living inside them are reading them. There is very little empirical work on how early-stage founders, investors, and ecosystem actors are actually interpreting artificial general intelligence as it arrives. Obschonka and Audretsch (2020) spell this out as an open call for field research, arguing that AI is already reshaping the conditions of entrepreneurship faster than the literature has caught up. The academic side of the conversation is rich on theory and thin on contemporary voices, while the industry side is loud on voices and short on method. This dissertation tries to occupy the space in between.

A note on register, because the next paragraph shifts tone. I wrote the opening in the first person because the question is personal. The argument for why the question matters structurally needs a different voice for a moment.

The structural stakes. AI capability is becoming an infrastructure layer. Cloud compute did this in the 2010s, TCP/IP did it in the 1990s, and each time the same pattern repeats: the thing that was once a differentiator becomes a substrate, and the competitive action moves one layer up. Epoch AI (2025) reports that five hyperscalers now control more than 67% of global AI compute. The Anthropic Economic Index (2026b) shows augmentation use at roughly half of Claude.ai traffic, with per-capita usage converging across US states even as capital concentrates in San Francisco. METR's (2025) task-horizon work finds the length of tasks AI systems can complete autonomously is doubling every four to seven months. That is the infrastructure picture. Sitting on top of it is an early-stage venture ecosystem that has not yet rewritten its playbook. This is the commoditisation problem: when the enabler is universally available, the rules for who wins and how change underneath the founders and investors making bets every day. Reading that change in real time is what I set out to do.

2.3 1.3 Where the study sits in the conversation

The dissertation inherits and extends three theoretical conversations.

The first is Davidsson, Recker and von Briel's (2020) external enablement framework, which reads environmental and technological shifts as the mechanisms that open, close or reshape the opportunity space available to new ventures. AI is a textbook external enabler. My data suggests, however, that the framework needs an amendment for enablers that become universally accessible almost instantly. When every founder can run the same code generation, the same retrieval pipeline, the same agent scaffolding, the enabler stops conferring advantage and starts acting as a threshold capability. That distinction did not exist in the 2020 paper because the enablers it described did not diffuse that fast.

The second is Bresnahan and Trajtenberg's (1995) general-purpose technology framework, which argued that steam, electricity and computing became engines of growth because they diffused across sectors and spawned complementary innovation. Historically these cycles took decades. Steam took about sixty years, electricity about thirty, computing about fifteen. My participants were describing cycles compressed to single product lifecycles. The GPT frame still applies, but the clock has been rewound.

The third is van Gelderen, Wiklund and McMullen's (2021) Delphi on the tension between everyday entrepreneurship and winner-take-most dynamics. Most of the literature treats this as a binary question: does AI democratise entry, or does it concentrate advantage? The interviews I conducted and the member-check conversations that followed in April 2026 pointed at a different answer. Both are happening at once, driven by the same technological force, and the right object of study is the paradox itself rather than either side of it.

2.4 1.4 Research question and objectives

The research question, updated to address the temporal anchor my advisor flagged on my first deliverable, is:

How will Artificial General Intelligence reshape the early-stage startup landscape over the 2025-2030 horizon that defines early-stage venture timing?

The 2025-2030 frame is deliberate. It matches the timeline on which seed-stage bets actually resolve, which is also the window over which most of my investor participants held AGI

capability estimates. A longer horizon would have drifted into speculation, a shorter one would have ignored the planning horizon that founders and investors use in practice.

Four research objectives organise the work:

1. To explore the expected changes in startup formation and operations resulting from the integration of advanced AI and AGI.
2. To investigate how AGI is anticipated to influence venture capital investment decisions and criteria at the seed and early stage.
3. To identify potential opportunities and threats that AGI presents to early-stage startups, from the perspective of startup founders and ecosystem experts.
4. To evaluate what capabilities and strategies entrepreneurs and early-stage support organisations will need to thrive in an AGI-augmented startup ecosystem.

2.5 1.5 How I studied this

A short overview here, the full account is in Chapter 3. I conducted ten semi-structured interviews in the San Francisco Bay Area between October 2025 and January 2026, with four founders, three investors (two VCs plus one angel), and three ecosystem participants: a programme director at a top-tier accelerator, an AI researcher who consults to early-stage ventures, and a community manager close to the day-to-day texture of the founder mood. Recruitment was purposive and snowball, routed primarily through the UC Berkeley network. Interviews ran between thirty-five and ninety minutes, nine were audio-recorded and transcribed with Otter.ai and manually corrected, and one participant declined recording. I analysed the data using Braun and Clarke's (2006) six-phase reflexive thematic analysis in NVivo, generating 127 initial codes and reaching saturation around the seventh interview. In April 2026 I returned to three participants for member checking, which turned out to be less a validation exercise than a longitudinal second round. The gap between the fieldwork and the member check was four to five months, and in a field moving this fast, that window produced meaningful new data rather than simple confirmation. I framed trustworthiness using Lincoln and Guba's (1985) four criteria (credibility, transferability, dependability, confirmability), which I discuss in detail in Chapter 3.

2.6 1.6 Contribution and roadmap

The central theoretical contribution I develop across the dissertation is what I call the **Bifurcated Ecosystem Model: A Framework for External Enabler Saturation**. It argues

that the AI-era early-stage ecosystem is splitting into a broad base of AI-consumer ventures built on commodity capability and a narrow tier of AI-producer ventures with proprietary data, custom models, or compounding domain advantage. The two layers are produced by the same force, not by opposite ones. Chapter 5 formalises the model into six testable propositions, each paired with a falsifiability condition, to move the contribution from description toward the kind of claim a future quantitative study could actually attack.

I did not expect this shape when I began. I thought I would find either a democratisation story or a concentration story, and I spent several weeks in late 2025 trying to force the data into one of the two. The shift happened when I stopped treating the tension as a contradiction to resolve and started treating it as the finding.

The rest of the dissertation unfolds as follows. Chapter 2 reviews the literature on AI as a general-purpose technology and external enabler, resource-based theory under commoditisation, entrepreneurial judgment under Knightian uncertainty, venture capital decision-making under hype, and the democratisation-concentration tension. Chapter 3 sets out the methodology in detail, including the Lincoln and Guba trustworthiness section, the two-part positionality statement that reflects my double role as student-researcher and active VC analyst, and an honest account of the data-quality issues that arose. Chapter 4 presents the findings as four themes and three counterintuitive patterns. Chapter 5 is the discussion, including the formalisation of the Bifurcated Ecosystem Model and a dedicated section on the AGI-to-ASI fragmentation that emerged during the writing period. Chapter 6 translates the findings into practical implications for three audiences: founders, investors, and accelerator or ecosystem institutions. Chapter 7 concludes with limitations and future research, including the frontier I am beginning to work on in my own research, the economic agency of autonomous AI agents themselves.

One last note before Chapter 2. I started this dissertation wanting to know what the early-stage world would look like by the time I fully entered it. Two years in, I am already inside it, writing IC memos on exactly the kind of deals my participants were pitching. That is lucky, and it is also disorienting, and it is why the question in the title is written in the future tense even though parts of the answer are already present.

3 Chapter 2: Literature Review

In this chapter I set out the conversations the dissertation is entering. I organise the literature around five themes: AI and AGI as general-purpose technology and external enabler (2.1); resource-based theory under commoditisation (2.2); entrepreneurial judgment and Knightian uncertainty (2.3); venture capital under uncertainty, herding, and AI hype (2.4); and the democratisation-versus-concentration tension (2.5). The chapter closes with a gap statement that motivates the empirical work in Chapters 3 to 7.

A note on register before the argument starts. This chapter leans analytical. I read the literature the way I read pitch decks in my current job: for the load-bearing claim, the numbers that support it, and the places where the argument is thinner than the confidence around it. I have tried not to paraphrase my way into fluency. Every subsection ends with a short paragraph in my own voice that links the theoretical material to the empirical work in Chapters 3 and 4.

3.1 2.1 AI and AGI as general-purpose technology and external enabler

3.1.1 2.1.1 The GPT frame and its historical clock

Bresnahan and Trajtenberg (1995) introduced the general-purpose technology (GPT) framework to explain why a small number of technologies produce outsized, economy-wide growth. For Bresnahan and Trajtenberg, a technology earns the GPT label only by clearing three tests at once: it has to spread broadly across sectors, it has to keep improving in capability, and it has to trigger waves of complementary innovation in the industries that adopt it. Steam, electricity, and information technology are the canonical examples. Each reshaped the organisation of production and the structure of the firm. Each eventually changed the conditions under which new firms could be founded.

The GPT frame is also a clock. Steam took roughly sixty years to diffuse through the British and then the global economy. Electricity took about thirty. Digital computing took about fifteen. Each cycle was shorter than the last, a pattern David (1990) describes as the “dynamo paradox”: productivity gains lagged adoption by decades, then accelerated as organisational forms caught up.

AI, taken as a GPT, compresses that clock further. Epoch AI (2025) reports that training compute for frontier models is growing at roughly 4.5x per year, with cost-adjusted capability growing at 3.5x per year and overall capability acceleration running about twice as fast over

the last two years as over the preceding decade. METR's (2025) task-horizon work, which measures the length of tasks AI systems can complete autonomously, finds a doubling time of four to seven months. Stanford HAI's (2026) AI Index reports that top frontier models now exceed 50% on Humanity's Last Exam, a benchmark that did not meaningfully exist two years earlier. The 2025 edition of the same report documented \$109.1 billion in US private AI investment; the 2026 edition updated that figure to \$285.9 billion for a single year.

The point is not that AI is unprecedented as a GPT. It is that the speed at which AI is diffusing leaves less time for the organisational and institutional catch-up that earlier GPTs got for free. Bresnahan and Trajtenberg's framework still applies; the clock has been rewound.

3.1.2 2.1.2 Digital entrepreneurship and the external enablement turn

The literature on digital entrepreneurship has tracked this acceleration from the demand side. Nambisan (2017) argued that digital technologies have dissolved traditional boundaries between ventures and their environments, turning entrepreneurship into a distributed, platform-mediated activity. Nambisan, Wright and Feldman (2019) extended this into a programme for digital entrepreneurship research, emphasising the role of infrastructure and ecosystem orchestration. Obschonka and Audretsch (2020) folded AI specifically into the conversation, arguing that machine learning is changing both the content of entrepreneurial action and the methods available to researchers studying it.

The most relevant theoretical move for this dissertation is Davidsson, Recker and von Briel's (2020) external enablement framework. The framework treats environmental and technological changes as enablers that alter the feasibility and profitability of new venture creation. An enabler is external to any given venture but shapes the opportunity space within which ventures are attempted. AI fits the framework cleanly: it lowers the marginal cost of building, it opens new product categories, and it changes the unit economics of early-stage ventures in measurable ways.

Davidsson et al. (2020) were careful to note that enablers do not act deterministically. A given enabler produces different outcomes depending on the mechanisms through which it is appropriated by founders and the agency of those founders in interpreting it. The framework is, in that sense, a language for describing how environmental change becomes entrepreneurial action. It was not written, however, for enablers that diffuse at the speed of a frontier model release and reach global availability within quarters rather than years.

Aghion, Jones and Jones (2019) frame AI more bluntly as an economic force, modelling it as a shift in the production function for ideas themselves. Their argument is that if AI begins to substitute for human research labour, the growth rate of ideas (and therefore of productivity) can change discontinuously. Agrawal, Gans and Goldfarb (2019) develop a complementary view, framing AI as a reduction in the cost of prediction, with first-order consequences for which economic activities remain valuable once prediction is cheap. Both papers are macro-level, but their implication for early-stage entrepreneurship is specific: the bottlenecks that previously determined which ventures were feasible have moved.

3.1.3 2.1.3 The AGI definitions problem

The academic literature handles AGI uncomfortably, and that discomfort is itself data. Bubeck et al. (2023), in the widely-discussed “Sparks of AGI” paper, argued that early GPT-4 exhibited behaviours consistent with partial general intelligence, while explicitly declining to claim full AGI. Morris et al. (2023), writing from Google DeepMind, offered the most serious attempt to operationalise the concept since Legg and Hutter (2007), proposing a two-dimensional “Levels of AGI” framework (narrow versus general capability on one axis, emerging to superhuman performance on the other) and explicitly rejecting single-threshold definitions. Aschenbrenner’s (2024) *Situational Awareness* treats AGI as a near-term planning assumption for national-security-level forecasting and estimates a 2027 arrival. Huang’s (March 2026) statement on Lex Fridman that “AGI is already here” in a functional sense sits at one end of a live disagreement; LeCun’s (2025-2026) repeated argument that current large language models will never reach AGI regardless of scale sits at the other. Amodei’s Davos 2026 remarks locate AGI within two years and forecast fifty per cent white-collar disruption within one to five years; Hassabis (2025) offers a five-to-ten-year horizon.

There is no settled academic definition of AGI. Bengio et al.’s (2026) International AI Safety Report acknowledges this explicitly and declines to pick one. The definitional disagreement is not a technical bug in the literature; it is a substantive feature of the phenomenon. AGI is a moving target whose perceived arrival correlates with the observer’s priors about which capabilities matter. What that means for the dissertation is that my participants brought the same definitional problem to the conversation and resolved it the same way: by substituting their working definition for the absent technical one.

3.1.4 2.1.4 Diffusion speed and the saturation condition

The acceleration claim deserves one more turn. Rogers' (1962, 2003) diffusion-of-innovations framework assumes a logistic curve whose slope is set by the interaction between innovation characteristics (relative advantage, compatibility, observability) and social-system variables. Historically, even highly advantaged technologies took decades to cross Rogers' adopter categories: electricity needed roughly forty years from the Pearl Street Station opening in 1882 to reach household saturation, and the personal computer needed about twenty-five years to reach two-thirds of US households. ChatGPT did the equivalent step in two months. The frontier-lab numbers are also tilted away from the consumer surface most adoption work has been calibrated to: OpenAI, Anthropic, and Google collectively report that their most-used surfaces are now enterprise APIs whose adoption is not publicly tracked on the same basis. The unit of comparison is not the same any more.

The diffusion literature does not have a clean frame for this. Geroski (2000) and Comin and Hobijn (2010) provide historical baselines for cross-country technology diffusion; both assume multi-decade timescales. Brynjolfsson, Rock and Syverson (2021) argue that general-purpose technologies follow a "productivity J-curve" in which measured output gains lag adoption by a decade or more, but the empirical basis of their argument predates the current model generation. Stanford HAI (2026) estimates that generative AI reached fifty-three per cent of US organisational adoption within three years of the first public release, which is the fastest documented diffusion of any GPT in the data. If we take the figure at face value, AI is crossing adoption thresholds on a timescale that the existing diffusion literature was not designed to describe.

This matters for the external enablement framing. Davidsson et al. (2020) assumed an enabler would remain a differentiator for long enough that the ventures appropriating it could build defensible positions around its use. When an enabler reaches majority adoption inside three years, that assumption fails. The theoretical move this dissertation makes in Chapter 5, extending the framework to the saturation case, is the direct consequence of the diffusion-speed anomaly described here.

3.1.5 2.1.5 Closing reflection

Before the interviews, I assumed "AGI" had a stable enough meaning in practitioner conversations to be a useful anchor. It does not. What I found, and what the literature above

supports, is that the term behaves less like a technical construct and more like a Rorschach blot on which founders and investors project their strategic assumptions. That shift from assumed stability to observed instability was the first move the literature made me make.

3.2 2.2 Resource-based view, moats, and the commoditisation problem

3.2.1 2.2.1 The VRIN framework at the limit

Barney's (1991) resource-based view (RBV) is the canonical framework for why some firms sustain advantage and others do not. The argument is that durable advantage requires the underlying resource to clear four hurdles at once: it has to be valuable, it has to be rare, it has to resist imitation, and substitutes have to be poor enough that competitors cannot work around it. The four criteria are usually shorthand as VRIN. The framework is descriptive: a resource that fails any one criterion cannot anchor durable competitive advantage.

Apply VRIN to AI execution capability in 2026 and the framework predicts a problem. A founder who can assemble a retrieval pipeline, fine-tune a model, and ship an agentic workflow is, in April 2026, doing work that is no longer rare and not especially hard to imitate. Anthropic's Economic Index (2026a) reports a twelvefold speed-up on college-level tasks under instruction; McElheran et al. (2024) document that large firms have adopted AI faster than small ones, which means the incumbents against whom a startup would compete for enterprise contracts already have the same capabilities. Stanford HAI (2026) reports GenAI adoption reaching fifty-three per cent of US organisations within three years of launch. If execution capability is universally accessible, it fails the inimitability criterion almost by definition.

The RBV does not predict that AI confers no advantage at all. It predicts that the advantage has migrated to the layer above execution: proprietary data, relationships, distribution, and judgment about what to build. Babina et al. (2024) confirm this migration empirically, finding that AI investments predict firm growth largely through product innovation rather than through operational efficiency, with the productivity benefits concentrated in firms that had complementary data assets before AI adoption. Sorgner (2024) reports a similar pattern in her review of AI and entrepreneurship: the firms benefiting most from AI are those with access to the complementary resources that AI does not provide by default.

3.2.2 2.2.2 Platform ecosystems and the tenant problem

Kretschmer et al. (2022) reframe the problem at the ecosystem level. Their argument is that platform ecosystems function as meta-organisations: loose collectives in which the platform owner sets the rules and the tenants operate within them. For AI-era startups, the platform owners are the frontier labs (OpenAI, Anthropic, Google DeepMind, xAI) and the hyperscalers that host them (AWS, Azure, GCP, CoreWeave). Epoch AI (2025) reports that five hyperscalers now control more than 67% of global AI compute. That concentration is the infrastructure floor on which every AI-dependent venture stands.

The implication of the platform ecosystems framework, read against the empirical concentration data, is that AI-dependent startups are tenants in a landlord economy they do not control. A pricing change from Anthropic, an API deprecation from OpenAI, a model update from Google: any of these can rewrite the unit economics of a venture overnight. Jacobides, Cennamo and Gawer (2018) describe the same dynamic at the industry level, arguing that ecosystem structure, not just firm strategy, determines value capture. The more recent literature on algorithmic management and platform power (Gregory et al. 2021; Cutolo and Kenney 2020) extends this to cases where the platform itself becomes an active agent in tenant operations, which is precisely what an AI model that sits inside a startup's product stack is.

3.2.3 2.2.3 Complementary assets and dynamic capabilities

Teece's (1986) complementary assets framework, developed to explain why the innovating firm often fails to capture the rents from its own innovation, offers a useful intermediate step between Barney's static VRIN picture and the migration dynamics above. Teece argued that whoever controls the specialised complementary assets (distribution, manufacturing, regulatory relationships, customer data) captures most of the value, even when the core innovation is contributed by someone else. Applied to the AI stack in 2026, the prediction is immediate: the frontier-model labs contribute the core innovation, but the durable rents accrue to whoever controls the complementary assets that the models depend on. Proprietary data, enterprise distribution, regulated customer relationships, and captive workflows are the complements. The RBV and the complementary-assets framework converge on the same answer from different directions.

Teece, Pisano and Shuen's (1997) dynamic capabilities extension, further developed by Helfat and Peteraf (2003) and Helfat and Raubitschek (2018) for the innovation-ecosystem case, adds the missing temporal dimension. In rapidly reconfiguring environments, advantage depends less on owning the right resources at a moment in time than on the capacity to sense, seize, and reconfigure around changing conditions. The framework has been criticised (Arend and Bromiley 2009) for tautology, and the criticism has weight. But the empirical observation the dynamic-capabilities literature captures is one that my participants described in their own words: the founders who were surviving AI's compression were those whose internal processes let them rebuild the product against the new baseline every few months.

3.2.4 2.2.4 Competence-destroying discontinuities, compressed

Tushman and Anderson's (1986) distinction between competence-enhancing and competence-destroying technological discontinuities is the third pillar of this subsection. Competence-destroying changes render prior expertise obsolete; competence-enhancing changes build on it. The original authors studied such discontinuities across cement, glass, and minicomputer industries, where the cycles were long enough for firms to reorganise around them.

The compression of the GPT clock discussed in Section 2.1 pushes Tushman and Anderson's framework to a limit. When competence-destroying discontinuities arrive at the cadence of quarterly model releases, the organisational response time they assumed is gone. A founder who is confident about a moat in January can watch it evaporate in March. The empirical reports confirm this: Faros AI's (2025) productivity paradox data show experienced teams absorbing new AI tools with non-trivial costs to their existing workflows, and METR's (2025) randomised controlled trial on experienced open-source developers found that AI tools actually slowed participants by 19% on complex tasks, a complicating result that the rest of this dissertation takes seriously.

3.2.5 2.2.5 Closing reflection

This framing helped me make sense of what P1 said when I asked about his competitive advantage five years ago. His answer was that advantage used to mean technical execution and customer relationships, and now it means proprietary data and judgment, and he was not sure what it would mean in three years. Read through Barney and Kretschmer together, P1 was not being vague; he was describing a moving VRIN perimeter inside a tenant economy. That reading is the theoretical scaffolding for what Chapter 4 calls the Resource Pivot.

3.3 2.3 Entrepreneurial judgment and Knightian uncertainty

3.3.1 2.3.1 Knight, McMullen-Shepherd, and the uncertainty core

Knight's (1921) distinction between risk and uncertainty remains the most useful single frame in the entrepreneurship literature. For Knight, risk and uncertainty are different beasts. Risk is the world in which you do not know the outcome but you can put numbers on the probabilities. Uncertainty is the world in which you cannot even price the probabilities. Knight argued that profit above normal returns is compensation for bearing uncertainty, not risk. If a future state is priceable, the market will price it; if it is not, entrepreneurial judgment is the mechanism by which capital is allocated under fundamental ignorance.

McMullen and Shepherd (2006) made uncertainty the theoretical heart of entrepreneurial action, defining entrepreneurial behaviour as action taken under conditions where means-ends frameworks are incomplete. Their argument has structured most of the subsequent literature on judgment and cognition in entrepreneurship.

3.3.2 2.3.2 The Townsend programme and computability

Townsend and Hunt (2019) brought Knight forward into the AI age directly. Their argument is that machine learning narrows the class of problems that remain genuinely Knightian by converting previously-unpriceable situations into priceable ones. The question they open, rather than close, is whether AI eliminates Knightian uncertainty or merely displaces it.

Townsend, Hunt and Rady (2024) push this further into the LLM era, arguing that large language models encounter a hard limit at the edge of reasoning about situations that do not yet exist. Some uncertainties remain fundamentally non-computable because they depend on reflexive social dynamics, institutional evolution, or future technological choices that have not yet been made. Those uncertainties cannot be reduced by more data; they can only be acted upon under judgment. The implication for AI entrepreneurship is direct: AI can compress the space of priceable problems, but it cannot eliminate the judgment layer that Knight identified a century ago.

Shepherd and Majchrzak (2022) offer a complementary framework for AI-augmented entrepreneurial cognition, arguing that the question is not whether AI replaces human judgment but how human and machine cognition combine in entrepreneurial decision processes. Their work is mostly theoretical; the empirical base for their claims is still thin,

which is part of what Obschonka and Audretsch (2020) called for in their field-research appeal.

3.3.3 2.3.3 Effectuation, discovery, and the cost-of-experimentation shift

Sarasvathy's (2001) effectuation framework proposes that expert entrepreneurs, under uncertainty, invert the classical causal planning logic: rather than starting from goals and pursuing the means to achieve them, they start from the means available and allow the goals to emerge through experimentation and commitment with stakeholders. The framework has been extended and tested across two decades of subsequent work (Read et al. 2009; Sarasvathy and Dew 2005; Arend, Sarooghi and Burkemper 2015) and remains one of the most-cited entrepreneurial-cognition lenses.

The relevance to AI is concrete. Effectuation assumes that experimentation is possible because the cost of a small experiment is low relative to the entrepreneur's resources. When AI collapses the cost of the first experiment by an order of magnitude (Anthropic 2026a reports a twelvefold speed-up on college-level tasks), the effectual mode should become more accessible, not less. The counterintuitive possibility, which several of my interviews surfaced, is the opposite: when experiments are cheap, the scarce resource becomes the willingness to commit to any one of them. Effectuation's "affordable loss" principle assumes that loss aversion acts as a natural filter; when loss is near zero per experiment, the filter is gone. This is the discipline problem that P1 named in his April 2026 member check, and it sits in a part of the literature that is not fully worked out yet.

Alvarez and Barney's (2007) discovery-versus-creation distinction adds a different angle. Their argument is that entrepreneurial opportunities either exist independently and are discovered, or they are socially constructed through action and interpretation. AI blurs the distinction. A founder who prompts a model for a business idea is neither strictly discovering nor strictly creating; they are performing a hybrid activity the framework does not cleanly describe. Packard (2017) and Alvarez, Barney and Anderson (2013) push the debate further without resolving it. My empirical data do not settle the question either, but they make clear that the discovery-creation axis is strained by AI-mediated opportunity identification in ways the literature will need to catch up with.

3.3.4 2.3.4 Judgment, skill, and AI-augmented action

The most useful frame for what AI does to a single founder's action set is Shepherd and Majchrzak's (2022) AI-augmented entrepreneurial cognition lens, introduced briefly in 2.3.2 and worth picking up again here. Their argument is that AI does not replace founder judgment; it changes the cognitive division of labour between founder and machine. What expands is the credible action set open to one person, because the machine takes on a wider share of the routine cognitive work. What does not change is the requirement that someone make the call.

That cleaner picture deserves one complication. Boussioux et al. (2024), working from a creative problem-solving experiment, found that generative AI widens the spread of ideas produced but narrows the quality distribution around them. The augmentation is real, but the trajectory is not monotonic toward better outcomes; it is toward more outcomes whose quality distribution may be tightening at the centre. The Resource Pivot framing in Chapter 4, where taste and discipline emerge as the new scarce resources, reads as a practitioner-side response to exactly that distribution.

3.3.5 2.3.5 Uncertainty migration as the load-bearing claim

Synthesising the literature above, the load-bearing theoretical claim for this dissertation is that AI does not reduce uncertainty in the Knightian sense. It moves it. Execution uncertainty (can we build this?) falls as AI handles more of the build. Model-dependency uncertainty (will our underlying API change underneath us?) rises. Regulatory uncertainty (what will the EU AI Act and its US analogues force us to do?) rises. Durability-of-advantage uncertainty (will the moat we think we have survive the next model release?) rises. The aggregate level of uncertainty a founder faces may not fall at all. It may even rise. What changes is where the uncertainty sits.

This is the uncertainty migration thesis. It is not my invention; it is the direction in which Townsend, Hunt and Rady (2024) point. My empirical contribution in Chapter 4 is to document the migration as it appears in practitioner interpretations.

3.3.6 2.3.6 Closing reflection

That was the thought that kept resurfacing in the interviews before I had a name for it. When P1 described his company's risk profile as having changed rather than shrunk, he was describing uncertainty migration in plain language. When P9 talked about hallucinations and

brittleness, he was identifying new uncertainties that had not existed in his earlier work. The literature gives the phenomenon a name; the interviews put a texture on it.

3.4 2.4 Venture capital under uncertainty: herding, horizons, and AI hype

3.4.1 2.4.1 Reputational herding and why it is rational

When P7, a VC associate at a prominent Bay Area fund, told me on tape that his fund's internal view was "AGI by 2030" and that "if we're wrong, we're wrong alongside everyone else," he was articulating the herding incentive in the first person. The literature has a name for what he described. Scharfstein and Stein (1990) is where the conversation on reputational herding starts: agents whose compensation is judged against peer performance have a structural reason to follow the crowd, because being wrong alongside a credible peer carries less reputational cost than being wrong alone. In venture capital, where fund performance is benchmarked against peers and where a portfolio bet that fails inside a popular thesis is more forgivable than one that fails outside it, the conditions for reputational herding are unusually strong.

Read against P7's line, herding is not irrationality. It is the rational answer to the incentives the fund structure creates. Sahlman (1990) describes the structural features of US VC that produce those incentives. Gompers and Lerner (2000) document how capital inflows into private equity inflate valuations in ways consistent with herding behaviour rather than with rational Bayesian updating on fundamentals. Kaplan and Stromberg (2003, 2004) analyse the contract structures sophisticated investors use to manage uncertainty at the deal level, finding that they actively design around the same incentive problems reputational herding creates.

3.4.2 2.4.2 AI hype, valuation, and the funding-gap data

The current AI capital cycle sits at the extreme end of this literature. The Letter Two (2026) reports that AI VC reached \$211 billion globally in 2025, with the San Francisco Bay Area capturing \$126 billion, or roughly 60% of the global total. Crunchbase's (2026) year-end review puts the AI total slightly lower at \$202.3 billion but reaches the same conclusion on concentration, noting that AI captured close to 50% of all global venture funding in 2025 and that OpenAI and Anthropic alone absorbed 14% of all global venture investment that year. 81% of Bay Area startup capital in 2025 went to AI-labelled ventures. Stanford HAI's (2026) AI Index puts total US private AI investment at \$285.9 billion, a 23x lead over China. SF Examiner (2026) reports that in Q1 2026, the Bay Area accounted for twelve of the top twenty

US rounds and roughly 75% of national startup funding in the quarter. Those are herding-level concentrations.

The WIPO (2025) Global Innovation Index adds a structural observation: global VC deal values rose 7.7% in 2024 while deal counts fell 4.4% for the third consecutive year. Capital is thickening into fewer, larger deals. That pattern is consistent with both herding (investors crowd into a narrower set of perceived winners) and with a structural shift in what early-stage venture finance actually finances (mega-rounds at Series B and later, with fewer but larger seed rounds).

The Carta (2026) data on founder ownership are the most revealing single data point. Solo-founded startups rose from 23.7% of new formations in 2019 to 36.3% in H1 2025. Thirty per cent of US startups founded in 2024 were solo-led. But those solo founders received only 14.7% of priced-equity capital. The gap between the founding rate and the funding rate is the Carta gap, and it is the clearest available empirical evidence that democratisation at the founding stage and concentration at the funding stage can coexist in the same dataset.

3.4.3 2.4.3 Hot-market dynamics and the power-law

Nanda and Rhodes-Kropf's (2013, 2017) work on investment-cycle dynamics provides the clearest theoretical link between aggregate capital abundance and deal-level behaviour. Their argument is that in "hot" markets, where capital is abundant and entry is crowded, the option value of being in a winner justifies a lower bar for the marginal deal, so investors fund more experimental, higher-variance ventures than they would in colder cycles. The prediction has two faces at once: hot markets produce more breakthrough successes because more experimental shots get funded, and they produce more failures at higher valuations because the bar for entry has fallen. The current AI cycle has every characteristic of a Nanda-Rhodes-Kropf hot market.

The venture return distribution is, as Kaplan and Lerner (2010) document, power-law rather than Gaussian: a small number of fund-returning investments produce most of the returns, with the modal investment returning close to zero. Gompers et al.'s (2020) large-scale survey of VC decision-making, published in the *Journal of Financial Economics*, reports that investors themselves identify team and market as the dominant decision criteria, with technology and business model following. The survey evidence is consistent with reputational-herding models: investors who cannot confidently price the underlying

technology substitute signals (team pedigree, market size, peer co-investment) that are easier to defend in front of their LPs.

Applied to AI, the power-law plus hot-market structure predicts exactly the funding-gap data the Carta (2026) report captures. A large number of experimental solo-founded ventures are funded to small seed rounds; a small number of team-led, capital-intensive AI ventures absorb most of the priced-equity capital. The distributional shape of the funding data follows directly from the theory, without requiring any specifically AI-driven mechanism.

3.4.4 2.4.4 AI at the other end of the table

The VC decision function is itself being reshaped by AI. Ma, Huang and Zheng (2022) document early applications of AI to deal screening and portfolio construction. Boubaker, Marcos and Rjiba (2022) trace how entrepreneurial finance is incorporating AI into due diligence workflows. Industry practice is running ahead of academic catalogue: Irwin-Hunt (2025) reports that AI dominated VC deployment in 2024, and CNBC (March 2025) documents Y Combinator's statement that its Winter 2025 cohort grew roughly ten per cent per week in revenue, a pace that requires algorithmic diligence and triage simply to track. Pitchbook's June 2025 reporting on Y Combinator's Spring 2025 cohort notes that 67 of 144 startups, roughly 47%, were AI agent companies, a concentration that creates obvious herding risks.

3.4.5 2.4.5 Timeline misalignment

The sharpest tension in the VC literature, for the purposes of this dissertation, is the misalignment between fund lifecycle horizons and AGI technical timelines. A typical seed fund operates on a ten-year horizon with an investment window of three to four years; the follow-on capital that rewards early bets must be deployable inside that window. Amodei's Davos 2026 two-year AGI estimate, Hassabis's five-to-ten-year estimate, and LeCun's continued skepticism of LLM-based pathways to AGI span a horizon that is longer than the time a typical VC can hold an illiquid early-stage bet. The structural consequence is that investor AGI timelines converge toward the fund's own horizon, not toward the frontier researchers' best estimates. Nanda, Samila and Sorenson (2020) document analogous horizon-driven distortions in earlier VC cycles; the same mechanism is visible in my interview data.

3.4.6 2.4.6 Closing reflection

The investor side of my sample spoke very differently from the technical side, and it took me a while to see why the misalignment was not a bug but a feature of how capital is priced. Scharfstein and Stein give the name: herding is what the incentive structure rewards. Gompers and Lerner give the amplification: capital inflows inflate valuation in ways that crowd out slower due diligence. Carta gives the shape: solo founders rise in count, but not in capital. That triangle anchors the empirical discussion in Chapter 5.

3.5 2.5 The democratisation-versus-concentration tension

3.5.1 2.5.1 Everyday entrepreneurship meets winner-take-most

Welter, Baker and Audretsch (2017) and van Gelderen, Wiklund and McMullen (2021) articulate the tension that structures this subsection. Welter et al. argue for a broader definition of entrepreneurship that includes everyday, non-heroic venture creation, pushing back against the field's over-focus on high-growth outliers. Van Gelderen et al., in a Delphi exercise drawing on editorial-board members of the leading entrepreneurship journals, surface a disagreement the literature has not resolved: is the AI era pushing the field toward everyday entrepreneurship (more people, lower barriers, more experiments) or toward winner-take-most dynamics (fewer, bigger, more concentrated)?

The Delphi is explicit about the disagreement and declines to resolve it. That non-resolution is the opening this dissertation steps into.

3.5.2 2.5.2 The democratisation evidence

Several sources support the democratisation reading. Eiss (2025) documents the rise of solo founders in Europe and North America. Carta (2026), already cited, puts the US solo-founder share at 36.3% of new formations. Graham's (August 2025) CNBC comments, although sceptical of AI-first company construction, implicitly acknowledge that the rate of new-venture creation in AI has accelerated beyond any recent baseline. Stanford HAI (2026) reports that 78% of organisations now use AI in at least one function, and TechCrunch (March 2025) reports that 25% of a recent Y Combinator cohort was producing 95%-or-more AI-generated code.

The CEO-visionary side of the conversation reinforces the democratisation frame. Altman's (February 2025) *Three Observations* essay posits a scaling law for value and a ten-times-per-

year cost decline for AI; his 2024-2025 claims about one-person billion-dollar companies, although partially walked back in later commentary, set the rhetorical ceiling of the democratisation narrative. Khosla (September 2025), at TechCrunch Disrupt, predicted a billion-dollar-revenue company with fewer than ten employees within five years. Hoffman's (2025) *Superagency* argues that AI extends the agency of individuals in ways that restructure the firm.

3.5.3 2.5.3 The concentration evidence

The concentration reading is equally well-supported. Epoch AI's (2025) report on compute concentration puts five hyperscalers at more than 67% of global AI compute. The Letter Two (2026) puts 60% of global AI VC in the San Francisco Bay Area. SF Examiner (2026) puts 75% of US Q1 2026 startup funding in the same geography. IMF (2025) Working Paper WP/25/76, *The Global Impact of AI: Mind the Gap*, frames the concentration at the cross-country level: AI capability and deployment are clustering in a narrow set of frontier economies, with the gap between frontier and laggard widening rather than narrowing. Anthropic's (2026b) Economic Index Geography paper adds a useful complication: per-capita usage is converging across US states even as capital concentrates in SF, which means the concentration is financial and infrastructural, not necessarily behavioural.

McElheran et al. (2024) extend the concentration finding to the firm level, documenting that larger firms have adopted AI faster than smaller ones. The intuitive expectation that AI would level the playing field between incumbents and startups does not match the firm-level adoption data. Incumbents moved first.

3.5.4 2.5.4 Institutional isomorphism and the ethics silence

DiMaggio and Powell (1983) give the cleanest organisational-theory lens on both the herding observed in 2.4 and the ethics-silence pattern Chapter 4 will describe. Organisations facing uncertainty converge on similar structures and practices through three pressures: coercive (regulation and politics), mimetic (copying credible peers), and normative (the shared training and professional networks the field flows through). All three are visible in the San Francisco AI ecosystem. Coercive pressure from the EU AI Act (Regulation 2024/1689) and the partial US analogue in the 2025 AI Action Plan is producing convergence on safety documentation inside the frontier labs. Mimetic pressure shows up in the near-uniform pitch-deck grammar Pitchbook (2025) reported across the YC Spring 2025 cohort. Normative pressure flows

through Y Combinator, the accelerator networks, and the investor-founder circuits Saxenian (1994) documented in earlier generations of Silicon Valley.

The framework predicts conformity where imitation is rewarded and neglect where it is not. Ethical reflection in early-stage AI ventures sits on the neglected side: no investor signal, no accelerator selection criterion, and no customer-contract clause pushes founders to foreground ethics in a pitch. The result is structural silence rather than personal indifference. Ergene, Banerjee and Hoffman (2021) and Wry and Zhao (2018) make adjacent arguments inside the institutional theory of sustainable and socially-oriented entrepreneurship; both would predict the pattern my interviews surfaced.

3.5.5 2.5.5 The tension as framework, not either-or

The usual move in the literature is to choose a side. Proponents of the democratisation reading tend to cite founder-count data and cost-of-founding trends. Proponents of the concentration reading tend to cite compute, capital, and geographic concentration. The honest reading of the evidence is that both are true, simultaneously, and produced by the same underlying force. AI lowers the cost of starting a venture, which produces more founders at the base. AI raises the value of proprietary data, relationships, and compute, which produces more concentration at the top. The same enabler drives both movements.

That synthesis is not in the literature yet as a stated theoretical position. The van Gelderen et al. (2021) Delphi surfaces the tension without resolving it; Welter et al. (2017) lean toward the everyday side; Wurth, Stam and Spigel (2023) approach the ecosystem side through a different lens, focused on regional entrepreneurial ecosystems rather than on the democratisation-concentration axis. The Bifurcated Ecosystem Model I develop in Chapter 5 is the formal attempt to close the gap.

3.5.6 2.5.6 Closing reflection

I went into the interviews expecting one side of this tension to win. I came out believing the tension itself is the finding. Reading the Delphi again after the fieldwork, I realised van Gelderen et al. had already posed the question in close to the shape my interviews answered; the contribution of the dissertation is to supply the answer they declined to give.

3.6 2.6 Gap statement and transition

Three gaps in the literature motivate the empirical work in the remainder of this dissertation.

The first is empirical. The literature on AI and entrepreneurship remains theory-rich and field-thin. Obschonka and Audretsch (2020) state this as an open call for research on how founders and investors are interpreting AI as it arrives. Townsend, Hunt and Rady (2024) push in the same direction from the judgment-and-cognition side. The industry and policy reports (Stanford HAI, Anthropic Economic Index, IMF, Carta) provide numbers but not interpretations. What the field lacks is a set of grounded accounts from practitioners who are making decisions at the frontier. This dissertation provides ten such accounts, triangulated across stakeholder groups and corroborated through a four-to-five-month member-check lag that itself becomes a second wave of fieldwork.

The second is theoretical. Davidsson, Recker and von Briel's (2020) external enablement framework was not written for enablers that become universally accessible within months of release. The framework correctly names the environmental mechanism but does not handle the saturation case in which an enabler confers no advantage because it is available to everyone. Extending the framework to the saturation case is the theoretical contribution I develop in Chapter 5 under the name of the Bifurcated Ecosystem Model. That model formalises the observation that universal access to an external enabler separates the ecosystem into a base that consumes the enabler and a tier that produces it, with the same underlying force generating both layers.

The third is typological. Practitioners report a founder-identity shift that the literature has not yet catalogued. The split between what I call Orchestrators (founders who conduct AI capability as a capability fabric) and Humanists (founders who retain the strategic-decision role and delegate execution to AI) appears across the interview data, is contested on the investor side, and is complicated rather than dissolved by P1's member-check resistance to the typology. No grounded typology of this shift currently exists in the published literature. Chapter 4 offers one.

The chapters that follow address the three gaps above in order. Chapter 3 sets out the methodology in detail, including the Lincoln and Guba (1985) trustworthiness framework. Chapter 4 presents the four themes and three counterintuitive patterns that emerged from the thematic analysis. Chapter 5 formalises the Bifurcated Ecosystem Model as six testable propositions, each paired with a falsifiability condition. Chapters 6 and 7 draw out the practical implications and the limits of the study.

One honest note before leaving the literature. The field is moving fast enough that some of the sources above were published between my original fieldwork in late 2025 and the writing of this chapter in April 2026. The Stanford HAI 2026 AI Index was released three days before I finalised this chapter. The Anthropic Economic Index Learning Curves report (Anthropic 2026c) was released in March 2026. I have treated this acceleration as a condition of the work rather than a problem to manage: the research question is about a moving target, and the literature about it is moving at similar speed. The discussion in Chapter 5 returns to this explicitly, because the AGI-to-ASI fragmentation visible in the Huang-versus-LeCun-versus-Amodei-versus-Hassabis spread during the writing period is itself part of the finding.

4 Chapter 3: Methodology

In this chapter I set out how I actually did the research, including the choices I made, the choices that were made for me by circumstance, and the places where the work could have been cleaner than it was. I have tried to write it the way I would explain it to a classmate, not the way a textbook would. The structure below follows Lincoln and Guba (1985) for trustworthiness, because that framework is both what the literature expects and what my advisor specifically asked me to foreground.

4.1 3.1 Philosophical positioning

The study is interpretivist, with a constructivist leaning. I take the view that the meaning of a phenomenon like “AGI” is not fixed by a technical definition and read off by interview participants. It is produced in the conversation between people who live inside that phenomenon. A founder who has watched three of his competitors die to the same model update has a different working definition of AGI than a VC partner who reads the term off a pitch deck. Both definitions are data.

This position follows Braun and Clarke’s (2019) defence of reflexive thematic analysis as a method that does not pretend to be a view from nowhere. The researcher is an instrument, the researcher has a perspective, and the honest thing to do is to make that perspective visible rather than to hide behind the passive voice. I write in the first person throughout the dissertation for this reason. It is not a stylistic indulgence, it is a methodological commitment.

What the position rules out is also worth stating. I did not approach this work as a hypothesis test. I did not have a quantitative instrument I was trying to validate. The field is not settled

enough for that yet. Before you can test a theory, you need to know what the practitioners are actually saying, and that is what this study tries to find out.

4.2 3.2 Research design

Given an exploratory research question on a fast-moving phenomenon with thin empirical literature, a qualitative, semi-structured interview design was the right fit. Semi-structured rather than structured, because I did not know in advance which threads each participant would pull hardest on. A founder running a legal-AI company and an angel investor writing checks twice a quarter would not have given useful answers to the same fixed question list.

I considered and rejected three alternatives. A survey would have reached more people but would have produced answers thin on the ambiguity that mattered, and the sampling frame for “people thinking carefully about AGI and early-stage startups in San Francisco” is not obvious enough to construct one from. A one-or-two-venture case study would have bought depth at the cost of cross-stakeholder triangulation. An ethnography would have matched the constructivist stance but required a year-plus SF residency I did not have.

The research question, as stated in Chapter 1, is:

How will Artificial General Intelligence reshape the early-stage startup landscape over the 2025-2030 horizon that defines early-stage venture timing?

The 2025-2030 horizon is not decorative. It sets the study at the window over which seed-stage bets actually resolve, which is also the window most of my investor participants worked with when they talked about AGI. That framing determined who I needed to talk to: people making decisions on that horizon, not people theorising about a century from now.

4.3 3.3 Sample and recruitment

I conducted ten semi-structured interviews with participants drawn from four stakeholder groups active in the San Francisco Bay Area early-stage ecosystem. The sampling strategy was purposive in the sense of Patton (2015), targeting participants likely to produce information-rich cases, with snowball sampling used to reach the harder-to-access investor and ecosystem segments after the first wave of founder interviews.

Table 1 below reproduces the participant matrix in the form it will also appear in Chapter 4 and in the appendices. All participants are anonymised, and identifying details of their companies, portfolios, and roles have been altered or removed to protect third parties.

Table 1. Participant Sampling Matrix

ID	Role	Age	Technical depth	AGI stance
P1	Solo Founder, Legal AI	29	High	Skeptical
P2	Co-Founder, Marketing AI	34	Medium	Cautious
P3	Solo Founder, EdTech AI	26	Very High	Optimistic
P4	Founder, HR SaaS	41	Medium	Pragmatic
P5	VC Partner	38	Low	Optimistic
P6	Angel Investor	45	Medium	Skeptical
P7	VC Associate	32	Low	Optimistic
P8	Accelerator Director	52	Low	Agnostic
P9	AI Researcher and Advisor	36	Very High	Realist
P10	Community Manager	29	Low	Observational

Four founders, three investors (two venture capitalists and one angel), and three ecosystem actors. The split was deliberate. I wanted enough founders to see variation in how AI-native and AI-adjacent ventures framed the same questions, and enough investors to test whether founder narratives matched the capital side or diverged from it. The third group, the ecosystem actors, ended up being more load-bearing than I initially expected: the accelerator director, P8, had visibility across more ventures than any single founder or investor and turned out to be the most useful participant for the counterintuitive patterns discussed in Chapter 4.

Recruitment was harder than I expected. I sent thirty-five cold LinkedIn messages to founders and investors who had been publicly quoted on AI and entrepreneurship. Exactly one converted into an interview. Every other participant came through a warm introduction, primarily via the UC Berkeley network I had built during my exchange at Haas, and secondarily through snowball referrals from P8, who acted as a gateway to three other ecosystem participants. The conversion ratio is itself a finding and a limitation. San Francisco is a warm-introduction economy, and a study of the San Francisco ecosystem that tried to bypass that cultural fact would have ended with three interviews and no paper.

I stopped at ten for two reasons. The first is analytical: I reached thematic saturation around interview seven and confirmed it at interview nine. The second is practical: I was running the

interviews in a narrow December-to-January window before returning to Europe for my next professional engagement, and past ten the logistics would have begun to compromise the quality of each individual session. Because saturation is a contested concept, as Braun and Clarke (2021) themselves have pointed out, I want to be specific about what saturation looked like in this dataset rather than only asserting it. By interview seven the Resource Pivot codes (taste, judgment as scarce, execution commoditisation) were stable: each new interview restated or refined codes already in the structure rather than introducing new ones. The Founder Identity codes (Orchestrator, Humanist, contextual judgment) were also stable by seven, though P3's resistance to clean classification kept the category boundaries productive rather than closed. The Timeline-Strategy codes and the three Counterintuitive Patterns were still moving slightly between interviews seven and nine: geographic-centralisation sharpened against P10's foot-traffic data at interview ten, and the ethics-silence pattern crystallised only retrospectively, after I noticed how few interviews had stayed with the topic. Braun and Clarke's critique of saturation is that the term often hides an absence rather than marks a presence; what I can show, rather than only assert, is that the dominant themes were stable by seven and that the moving codes are named explicitly. The April 2026 member-check responses then deepened existing structures rather than introducing new themes, which I take as the strongest available signal that the thematic structure was complete enough for this sample.

The warm-intro bias deserves its own acknowledgement. My sample over-represents people who were accessible to a Berkeley-affiliated student in late 2025 and under-represents people who are either too senior or too insular to be reached that way. I return to this in the limitations section of Chapter 7.

4.4 3.4 Interview conduct and data quality

Interviews ran between thirty-five and ninety minutes, with most falling in the fifty-to-sixty-five minute range. I used a semi-structured guide organised around three blocks: the participant's current work and how AI had changed it, their working definition of and timeline for AGI, and their views on how capital, talent, and strategy were likely to respond. The full guide is reproduced in Appendix A, with annotations on the probes I used most often and the ones I dropped after the first three interviews because they produced thin answers.

Nine of the ten interviews were audio-recorded with the participant's consent, transcribed with Otter.ai, and manually corrected against the original audio for accuracy. The tenth (P5, a

VC partner) declined recording. That is worth pausing on. Recording refusal is itself a signal in an environment where most technologists are comfortable on the record. In that case I wrote a narrative summary within two hours of the conversation while my memory was fresh, and I have flagged reconstructed passages in the transcript file so that no paraphrase is treated as a direct quotation in the findings chapter.

Two further data-quality issues need honest treatment.

First, the recording of P4's interview was corrupted at around the twenty-minute mark, and I did not notice until I sat down to transcribe it the following day. The remaining forty minutes of the conversation were reconstructed from handwritten notes I had taken during the interview, supplemented by voice-memo observations I had recorded on my phone on the walk home. I wrote the reconstruction within twenty-four hours of the interview, flagged reconstructed passages in the transcript, and when a reconstructed statement is quoted in Chapter 4 it is cited as paraphrase rather than direct quotation. The coding of the first twenty minutes, which was recorded cleanly, anchors the reconstruction.

Second, P7, a VC associate at a prominent Bay Area fund, was noticeably guarded throughout the conversation. Interviewing P7 was the most methodologically uncomfortable experience of the research process. His answers were careful, hedged, and at several points explicitly off-record. I read his hedging as institutional rather than personal because the dimensions he was guarded on (peer-following, internal AGI estimates) are exactly the ones a fund's communications policy shapes most tightly, but I cannot rule out that some of the reluctance was skepticism toward the research itself. I considered dropping the interview from the dataset and decided against it, because the institutional guardedness itself was data, and because his comments about following "credible peers" under uncertainty resonated with Scharfstein and Stein's (1990) reputational herding framework in a way I could not ignore. I triangulated P7's substantive claims against P5 (a VC partner at a different firm) and P6 (an angel investor), and I flag in the findings chapter where a P7 observation is corroborated and where it stands alone.

Three of the ten interviews therefore carry a visible caveat (P4 recording corruption, P5 recording refusal, P7 institutional guardedness). I treat that as a feature of qualitative fieldwork in a high-secrecy ecosystem rather than a failure of the study, and I flag in the findings which claims rest on clean data and which rest on reconstructions.

4.5 3.5 Analytical approach

Analysis followed Braun and Clarke's (2006) six-phase reflexive thematic analysis, conducted in NVivo 14. The phases as I worked through them were: familiarisation, initial code generation, theme searching, theme review, theme definition, and the production of the written analysis.

Familiarisation meant reading every transcript twice before coding anything, and listening to the audio a second time for the interviews I felt least sure about. The initial coding phase produced 127 codes, a mix of deductive codes derived from the literature (e.g. "external enabler", "uncertainty migration", "VRIN resource") and inductive codes that emerged from the data itself (e.g. "taste as scarce resource", "AGI moat", "ethics silence"). The deductive-inductive split was roughly forty-sixty in favour of inductive, which is consistent with the exploratory character of the research question.

Theme searching and review condensed the 127 codes into four major themes and three counterintuitive patterns, which I present in Chapter 4. One of the analytical pivots I want to flag here, because it matters for the dependability of the analysis, is the decision to retain the Orchestrator-versus-Humanist typology in the final framework even after P1 pushed back on it during member checking in April 2026. P1's view was that most good founders live between the two poles. That critique is fair, and I have folded it into the chapter as a refinement rather than a rejection. The typology stayed because the ambiguity P1 pointed to was itself analytically informative: it showed that the archetype distinction was real enough to orient the field and soft enough to misclassify specific founders. I return to this in Section 3.7 under dependability.

Saturation arrived around interview seven and was confirmed at interview nine. Analytical work continued past that point, through member checking and into the writing.

4.6 3.6 Member checking (April 2026)

In April 2026, four to five months after the end of the fieldwork window, I returned to three participants (P1, P3, P8) with a structured follow-up. The three were chosen on two criteria. First, cross-stakeholder coverage: P1 a founder, P3 a founder of a different technical depth and AGI stance, P8 the ecosystem participant whose accelerator vantage point gave him visibility across more ventures than any single founder or investor in the sample. Second, engagement signal at the original interview: each of the three had pushed back, qualified, or

extended their own claims in the original conversation, which suggested they were the most likely to engage substantively with a structured follow-up. I sent the same template to a fourth participant (an investor) but did not receive a response within the window, and I did not push further because the asymmetry of three founders or ecosystem actors and zero investors at the member-check layer is itself a transparent caveat that travels with the findings. Each of the three who did respond received a two-page summary of the themes as I had coded them, a short set of direct questions about interpretations that were specifically attributable to them, and an open invitation to correct, extend, or complicate the analysis. All three responded at length. The full responses are archived and excerpted in the findings and discussion chapters; full extracts sit in Appendix D.

I want to frame the member checking as a methodological choice rather than a courtesy. The conventional framing of member checking is as a validity exercise: you return to participants to confirm that your interpretation of their words is the interpretation they would recognise. That was part of what I did. But the four-to-five-month gap between the original interviews and the follow-up, in a field moving at the pace of late-2025-to-early-2026 AI development, turned the exercise into something closer to a longitudinal second wave. P8 had run a Winter 2026 accelerator cohort between the interview and the member check and produced fresh data on remote-team experiments, prototyping velocity, and curriculum redesign. P1 had been through an investor conversation that crystallised the “AGI moat” anecdote I quote in Chapter 4. P3 had rebuilt his product twice. This is not confirmation, it is enrichment. The original data remain the anchor; the member-check responses are treated as a second, lighter wave of fieldwork.

I note the asymmetry honestly: three of ten is not full coverage. The three who responded were the three with the highest engagement at the original interview, which means the member-check layer is itself subject to a selection effect. I treat it as corroboration rather than validation of the full sample and return to this in the limitations section of Chapter 7.

4.7 3.7 Trustworthiness (Lincoln and Guba 1985)

This section is deliberately long. My advisor’s feedback on the third deliverable was that validity and reliability needed to go beyond a one-paragraph mention of saturation, and Lincoln and Guba’s (1985) four criteria are the standard framework qualitative researchers use to meet that bar. I address each criterion explicitly.

4.7.1 Credibility

Credibility is the qualitative analogue to internal validity. It asks whether the account I am giving captures the meaning my participants actually made of their situation, or whether I have substituted my own. Three mechanisms support it in this study.

The first is member checking, detailed in Section 3.6. Three participants (P1, P3, P8) reviewed the thematic structure attributed to them and responded with corrections, extensions, and in one case a useful refinement of a typology.

The second is within-sample triangulation. The sample was constructed to produce cross-stakeholder tension, not single-voice consensus. Where an investor and a founder said the same thing, I treat the finding as well-supported. Where they disagreed (as they did on the Orchestrator versus Humanist question), I treat the disagreement itself as the finding, rather than forcing resolution toward one side. P7's guardedness, for example, is only interpretable against P5 and P6's more candid accounts.

The third is triangulation with contemporary secondary sources. Throughout the analysis I checked the patterns emerging from the interviews against quantitative evidence from the Stanford AI Index (2025, 2026), the Anthropic Economic Index (2026a), Carta (2026), the IMF Working Paper WP/25/76 on global AI inequality, and METR's (2025) task-horizon work. When the interview data diverged from the secondary data (for instance, on the pace at which AI adoption was spreading outside the Bay Area), I have tried to follow the divergence rather than paper over it.

4.7.2 Transferability

Transferability sits where external validity sits in the quantitative tradition, and Lincoln and Guba treat it as a question the reader answers rather than one the researcher claims. My task is to provide enough thick description of the context that a reader can decide whether the findings apply to the setting they care about.

The setting of this study is specific. San Francisco is not a typical startup ecosystem, it is the extreme end of one. The IMF's 2025 global AI report and The Letter Two (2026) both document the degree to which early-stage AI capital, talent, and compute are concentrated in the Bay Area. P8's accelerator data (roughly 60% SF-based in 2023, 85% in 2025) is a local signal for a trend visible at the national level. A reader wanting to transfer these findings to a less concentrated hub (Paris, London, Toronto, Bangalore) needs to read the findings with that

asymmetry in mind. I try to describe SF thickly enough that readers can do that work, rather than assuming the findings generalise.

4.7.3 Dependability

Dependability is about whether the analytic trail is reproducible enough that another researcher, given the same transcripts and coding decisions, would arrive at a recognisably similar set of themes.

Three elements support it. First, the coding structure is documented in full in Appendix G, from the 127 initial codes through the collapses and expansions that produced the four themes and three counterintuitive patterns. Second, analytical pivots are named rather than smoothed over. The decision to retain Orchestrator-Humanist despite P1's resistance, the decision to merge several execution-related codes into the single "Resource Pivot" theme, the decision to separate Geographic Centralisation out of the Timeline-Strategy nexus after interview seven: each is logged with a short rationale. Third, the theme evolution across the six phases of Braun and Clarke's framework is reproduced in Appendix G in table form, so that a reader can see how, for example, the inductive code "taste" grew into the theme of the Resource Pivot after interviews four through seven but was only joined by its twin "discipline" after P1's April 2026 member check.

4.7.4 Confirmability

Confirmability is the question of whether the themes emerged from the data or from my priors. The conventional tool is a reflexive journal, maintained contemporaneously through the fieldwork and coding period.

A note on the journal artefacts. The canonical model assumes a single, unified document maintained throughout the work. I did not maintain that document. What I did maintain was a partial set of journal-equivalent artefacts: post-interview fieldnotes written within an hour of each conversation, analytical memos written into the NVivo project during the coding phase, and voice memos recorded on most of the walks home from interviews. Read together, these constitute a partial reflexive trail, just not the unified journal the canonical model assumes. When I began writing the dissertation in early 2026 I needed to reconstruct some of the interpretive choices that the partial trail did not capture, and the reflexive journal excerpts reproduced in Appendix E are accordingly a mix of contemporaneous fieldnotes, contemporaneous analytical memos, and retrospective reconstructions written within a short

window after writing began. Retrospectively reconstructed passages are flagged as such. I prefer to name this limitation than to hide it. The positionality statement in the next section compensates by making my standpoint explicit, and the partial-trail artefacts compensate by anchoring the contemporaneous side of the reflexive work; together they substantially narrow the gap left by the absence of a unified journal, without fully replacing it.

4.8 3.8 Reflexivity and positionality

I occupy two positions in relation to this research simultaneously, and both need to be on the page.

The first I have previously called “pedagogical distance.” I am a twenty-one-year-old European student, temporarily in San Francisco on exchange at UC Berkeley, writing a dissertation for a French business school. In most interviews, I was younger than the participant by a decade or more. I had no standing, no portfolio, no fund. That outsider status had two effects. It made tacit knowledge more visible, because participants took the time to explain things they would have glossed over with a peer. And it changed the kind of questions I could ask, because a participant has different incentives when pushing back on a student than on a reporter or a rival. P7 agreed to the interview partly, I suspect, because a student posed low reputational risk; P9 likely simplified his technical explanations for the same reason. Both of those effects cut both ways: the access was real, and the simplification was a filter on the data.

The second position is what I have started to call “participant proximity.” Between the end of the fieldwork and the writing of this dissertation, my professional role changed. Since January 2026 I have been a VC analyst at Neo Founders and One Green in Paris, screening thirty to forty seed and Series A deals per month, writing investment-committee memos, and helping to close SPVs. As I noted in the Introduction, I am writing this dissertation as a participant in the thing I am studying rather than as an outside observer; this section makes the participation specific. On a smaller scale and in a different market, I am inside the workflow my investor participants described. This is the second of three deployments of that biographical fact in the dissertation, following the one in the Introduction and preceding the one in Chapter 6 on practical implications for investors.

The participant-proximity position has two effects on the analysis that I have tried to monitor rather than hide.

As a methodological asset, it functions as a continuous secondary exposure to the phenomena I am studying. When P1 described the “AGI moat” question in his April 2026 follow-up, I recognised the shape of the conversation because I had sat in similar rooms in Paris in the intervening weeks. When P5 and P6 disagreed about whether to back Orchestrator founders, I had seen both types of memo written in my own deal flow. This is not a substitute for the interview data, and I treat it strictly as corroboration rather than evidence, but it changes the credibility framing of the analysis. I am not an observer inferring what investors think. I am currently writing some of those investor documents myself.

As a bias, the same exposure predisposes me to see the four themes everywhere. The risk is obvious. A researcher who codes for the Resource Pivot on Monday and then reviews three seed-stage pitch decks on Tuesday will be primed to find the Resource Pivot in the decks regardless of whether it is really there. My mitigation has been to return to the original transcripts when a pattern began to feel too clean, and to treat the April 2026 member-check responses as a correction layer on the 2025-era interpretations. Where my current professional intuitions conflict with the interview data, I have tried to follow the data.

Two further biases are worth naming briefly. I am an AI enthusiast. I have built two on-chain protocols (UMIP in 2025 and AgentEscrow in 2026) that assume autonomous AI agents will become economic participants, and I hold that view publicly. That enthusiasm makes me more likely to take claims about AI capability at face value and less likely to probe skeptical positions (like P9’s “hallucinations and brittleness” argument) for longer than I did. I flag this and have tried to compensate by engaging the skeptical literature (Gownder 2026 at Forrester, METR 2025 on AI slowing experienced developers) seriously in the discussion. My second bias is that I entered the fieldwork having already written about AI as an infrastructure-layer shift. Coming in with a lens is not in itself a failure of method; failing to make the lens visible would be.

4.9 3.9 Ethics

The study sits inside ESSEC’s research ethics framework and was cleared at the D2 stage. Every participant agreed to take part on an informed basis: nine signed the written consent form reproduced in Appendix B before their interview, and P5 gave verbal consent that I wrote up in my notes within two hours of the conversation and signed off in writing at that point. What each participant signed up to was the interview itself, audio recording where they

accepted it, the later use of anonymised quotations in this dissertation and in any academic work that grows out of it, and an open right to pull out up to the date of submission.

Anonymisation runs at three distinct layers. At the person layer, everyone is referred to by a P-number allocated in the order they were interviewed. At the institutional layer, I have stripped or generalised company names, fund names, portfolio holdings, and accelerator branding throughout. At the detail layer, I have softened specific product lines, syndicate partners, and micro-niches to the point where a reader outside a participant's immediate circle should not be able to trace the interview back to them. P1, P3, and P8, the three I re-contacted in April 2026, were additionally offered a veto over any excerpt attributed to them; none used it.

Third parties need their own treatment in a sample this small. Several of my participants brought in people who had not agreed to be in a research dataset: co-founders, partners, operating peers, portfolio-company teams. Anything that could identify those third parties, by name or by distinctive detail, has been pulled out of the transcripts and out of the dissertation.

A log of the AI tools used during the research process is reproduced in Appendix F, in line with ESSEC's 2025 guidance on AI use in academic work.

One closing observation before Chapter 4. Writing this chapter honestly was harder than I expected. The temptation in a methodology chapter is to round off the corners so that the process looks cleaner than it was. I have tried to resist that. The recording corruption, the recording refusal, the institutional guardedness, the reflexive journal I did not keep, the selection effect in the member check: all of them are on the page above, because a reader who does not know about them cannot weight the findings properly. The findings, presented in the next chapter, rest on the design described here, with its real shape.

5 Chapter 4: Findings

In this chapter I present the findings from the thematic analysis described in Chapter 3. The material is organised around four themes, the fourth of which contains three counterintuitive patterns that I present as unexpected results. The chapter closes with a synthesis that shows how the themes connect and prepares the ground for the theoretical discussion in Chapter 5.

A note on framing before I start. I did not go into the fieldwork expecting to find a paradox. I expected one side of the democratisation-versus-concentration debate to win on the ground,

and I spent several weeks in late 2025 trying to force the data toward one side or the other. The shift in my thinking came when I stopped treating the tension as a contradiction to resolve and started treating it as the structure of the answer. The most concise version of that answer is the line P8 gave me in his first interview, which I have cited in earlier work and which has since become the organising frame of the dissertation: *“It’s never been easier to start a company, and never been harder to win.”* That line stayed with me because it refuses to sit on either side of the usual argument. Both halves are true at once, produced by the same technological movement, and the empirical chapters that follow are an attempt to describe why.

What I present below is a landscape (one of the few uses I allow that word in this chapter, because it sits in the dissertation title) in which the easy answers did not survive contact with the data. The AI-era early-stage ecosystem is not simplifying. It is splitting along axes the existing literature has named but not yet theorised together. The four themes are: a shift in what counts as a scarce resource, a shift in how founders understand themselves, a misalignment between the timelines that founders and investors work with, and three patterns that run against the expectations the literature would have given me in advance.

5.1 4.1 Theme 1: The Resource Pivot, from Execution to Judgment

The most consistent finding across the interviews, and the one that appeared earliest in the coding process, was that AI has rewritten what counts as a scarce resource in early-stage ventures. The first barrier in most 2010s founder stories was whether you could ship the thing at all; that barrier has thinned to the point of disappearing. Across the sample, the story my participants were telling was about a new equilibrium: writing the product is assumed, and the work that distinguishes a serious venture from the volume around it is the work of choosing which product to write.

P1, a legal-AI solo founder with a strong technical background, put it most directly. When I asked him what his competitive advantage had been five years earlier, he answered:

“Five years ago, my competitive advantage was that I could code. Now anyone can code with Claude or Cursor. So what’s my advantage now? It’s that I know what to build. The AI can write the code. It can’t tell you which problem is worth solving.”

The quote compresses the whole theme into four sentences. Coding used to be a VRIN resource in Barney's (1991) sense. It is not one any more. What has replaced it, at least in P1's articulation, is judgment about problem selection.

P8, the accelerator director, saw the same compression from the institutional vantage point. In the cohort he was running at the time of interview, most teams were standing up a working prototype inside the first two weeks, and the content of the mentorship hours had swung away from *"how to build"* and toward *"what to build and for whom"*. P5, a VC partner, used a phrase I flagged twice in my fieldnotes because it is a useful compression even when it is lossy: her firm now screens for *"taste-makers rather than engineers"*. P7, a VC associate at a prominent fund, said he was working through *"50 pitches per week starting with 'We use GPT to...' with no actual moat beyond prompt engineering"*. The vocabulary differs across founders, investors, and ecosystem actors, but the movement underneath it is the same: the scarce capability has migrated up the stack.

Underneath the shared pattern, the data surfaced a sharper division. Several participants were drawing a line between two distinct kinds of AI-era venture, and this is the point in the chapter where the descriptive account starts to become a theoretical one. P6, an angel investor with a long memory for hype cycles, offered two labels I retained in the coding structure: AI *"consumers"* and AI *"producers"*. Ventures on the consumer side sit on top of off-the-shelf frontier APIs and fight for share in categories with very little vertical differentiation. Ventures on the producer side own something that compounds over time: a proprietary dataset, a fine-tuned or custom model, a distribution channel that is hard to recreate. P9, the AI researcher and advisor, restated the same split as a data-layer argument: training-data control is where durable advantage now sits, while the weights and compute that used to be the scarce inputs are on a fast commoditisation track.

The consumer-producer division is the empirical basis for what I call in Chapter 5 the Bifurcated Ecosystem Model. The two layers are not separate ecosystems sitting side by side. They are produced by the same technological force. Cheap, universally available AI lowers the cost of starting a venture, which floods the base with consumers. The same cheap, universally available AI raises the marginal value of proprietary data, enterprise relationships, and regulated customer access, which strengthens the peak occupied by producers. Figure 1 summarises this as the graphical form the later chapters build on.




Figure 1. The Bifurcated Ecosystem Model: A Framework for External Enabler Saturation

Figure 1. The Bifurcated Ecosystem Model: A Framework for External Enabler Saturation. The same technological force (commoditised AI) enables the broad base of AI consumers while strengthening the narrow peak of AI producers. Central paradox after P8: *"Never easier to start, never harder to win."*

P6 translated the split into an investor diligence heuristic that I found useful enough to quote in full, because it captures the texture of how the pattern lands in practice:

"I ask founders to show me three decisions where they didn't use AI, even though they could have. If they say 'we use AI for everything!' that's a red flag. The best founders use AI aggressively but selectively."

The heuristic P6 described, which I came to code under *"selective restraint"*, kept surfacing in other interviews as the marker that separated sophisticated founders from the rest, often implicitly in what a founder chose not to automate rather than what they did.

The member check with P1 in April 2026 pushed the theme further, and I want to treat his correction explicitly because it is the kind of refinement the original coding did not fully

capture. P1’s follow-up was that “*taste*” is necessary but not sufficient. He added a second term: discipline. His argument was that AI has made building so cheap that the actual bottleneck is no longer capability, or even the ability to identify a worthwhile problem, but the ability to say no to features that are easy to build and wrong for the product. He told me he had watched three competitors in his space launch and fold since our first interview, all of them built on commodity APIs with no proprietary legal-workflow data. The pattern P6 named as a diligence heuristic, P1 had lived as a survival condition.

What struck me about the Resource Pivot after the member check is that it is not one pivot but two, stacked. The first is the commoditisation of execution, which the original interviews captured. The second is the emergence of discipline as the new constraint once building is cheap, which P1 surfaced only after four additional months of market exposure. I have not tried to smooth the two into a single line in the chapter because I think the stacking is itself informative. The scarce resource moves as fast as the enabler underneath it.

5.2 4.2 Theme 2: The Founder Identity Rewrite

The second theme followed from the first but carried a different charge. Once the tasks that used to define the founder’s contribution are absorbed by AI, the question of what a founder *is* stops being rhetorical. Several participants described something closer to an identity shift than a tooling shift. Two archetypes emerged cleanly enough in the coding that I kept them as a typology, while being honest about the places where the typology started to leak.

I call the two archetypes Orchestrators and Humanists. P3, a solo EdTech founder with very high technical depth, described himself as a “*conductor of an AI orchestra*”. P4, an older founder building an HR SaaS product, described himself as a “*strategic decision-maker*” whose core work was relational rather than technical. Those two phrases became the anchor quotes for the two ends of the typology. Table 2 lays out the contrasts as they appeared across the sample.

Table 2. Founder Identity Typology

Dimension	Orchestrators	Humanists
Self-concept	“Conductor of an AI orchestra” (P3)	“Strategic decision-maker” (P4)
Core activity	Managing AI systems,	Relationship-building,

Dimension	Orchestrators	Humanists
	optimising workflows	contextual judgment
View of AI	Instruments to coordinate and direct	Tools for efficiency, subordinate to human insight
Emotional tone	“Am I still an entrepreneur?” (multiple participants, observed by P10)	“The job hasn’t changed” (P4)

The typology is not only about self-concept. It shows up in how founders allocate their time, what they hire for, and what they optimise around. Orchestrators build tooling pipelines and spend their mornings in agent loops. Humanists spend their mornings on customer calls. Both can be successful; both are recognisable on the ground.

P10, the community manager whose daily work put her closer to the ambient founder mood than anyone else in the sample, described the psychological load of the shift:

“I hear it monthly now: ‘If AI can do my job, what am I good at?’ Some founders get energised by that question. Others get genuinely depressed.”

P10’s phrasing was the first thing that made me treat the theme as identity rather than tooling. The founders she was watching were not debugging, they were renegotiating.

P1 pushed back on the premise that AI could replace founder judgment at all, and the quote he gave me is a small masterpiece of what the Humanist side of the typology actually defends:

“I can hear pain points in customer calls that AI wouldn’t flag because I understand context in ways that don’t translate to text.”

That line is not a claim that AI is weak. It is a claim that judgment has an embodied, contextual quality that does not reduce to text, and therefore does not yet reduce to a model. Whether the claim survives the next generation of multimodal systems is an empirical question the dissertation cannot settle. What matters for the theme is that P1 was articulating, in his own language, a version of the Knightian uncertainty argument Townsend, Hunt and Rady (2024) describe at the theoretical level.

The typology held for most of the sample, with one exception that ended up being the most productive case in the chapter. P3 refused to sit cleanly inside either pole during the original interview. In the same half-hour he framed himself as an “*AI conductor*” and as someone

who “*hates when people stop thinking for themselves*”. On the strength of the rest of his behaviour, orchestration-heavy across the board, I coded him as an Orchestrator, while noting in my memos that the label was doing slightly less work on him than on the others. The ambiguity itself felt analytically significant.

Where the theme sharpens is on the capital side. The investor sub-sample is not aligned on which archetype to back. P5, the VC partner, writes cheques to founders who “*can orchestrate AI, humans, and capital*”, which lands as an unambiguous Orchestrator endorsement. P6, the angel investor, is looking the other way and wants “*founders less bought into AI hype who focus human energy on relationships and brand*”. Two investors at the same stage are writing different term sheets against the same kind of deck. The identity question is therefore not only a private matter for founders; it is a market-level signal about what the capital side believes the next ten years will actually reward, and that signal is currently incoherent.

The member checks added two refinements that I have kept in the final version of the typology rather than rewriting around. P1 wrote in his April 2026 response that he was “*neither*” archetype, that the framing implied a binary he did not recognise in practice, and that most of the good founders he knew sat somewhere in between and got into trouble when they committed too hard to either pole. I take this seriously as a correction of the typology’s usability, and I fold it into the presentation above by treating Orchestrator and Humanist as poles of an axis rather than as exclusive categories.

P3’s member check pushed the same correction from a different angle. His proposal was that the real distinction is not Orchestrator-versus-Humanist but “*building FOR AI vs building WITH AI*”. He classified himself as a Humanist who uses Orchestrator tactics, and added that the ambiguity in his original interview was probably the point. I have not rewritten the typology around P3’s proposed axis because my sample does not have enough cases to defend a second typology, but I flag the proposal in Chapter 5 as a direction for future quantitative validation. One of the six propositions I develop in Chapter 5.4 takes P3’s axis seriously as an alternative formulation.

The ambiguity at the edge of the typology is the finding, not a defect of it. A clean typology that every participant fit cleanly into would have meant the coding had forced the data; the places where the categories leak are where the theoretical work begins.

5.3 4.3 Theme 3: The Timeline-Strategy Nexus

The third theme connects participants' beliefs about AGI timelines to the strategic decisions they make now. The simplest version of the pattern is that technical depth predicts AGI skepticism, and investor timelines cluster around fund horizons rather than around the best available technical estimates. Table 3 gives the participant-level view that lets the argument be read off the data.

Table 3. Participant Perspectives Matrix

Participant	AGI Timeline	Primary Concern	AI Strategy	Stakeholder Group
P1	10+ years	Sustainable moats	Selective adoption	Founders
P2	8-10 years	Model dependency	Model-agnostic	Founders
P3	3-5 years	Missing the wave	AI-native, aggressive	Founders
P4	15+ years	Unit economics	AI as cost tool	Founders
P5	5-8 years	Portfolio positioning	Mixed approach	Investors
P6	10-15 years	Bubble risk	AI-resistant moats	Investors
P7	5-7 years	Deal flow quality	AI-first bias	Investors
P8	Agnostic	Founder quality	Tool-agnostic	Ecosystem
P9	15-20 years	Safety and alignment	Research-focused	Ecosystem
P10	Observational	Culture shift	Observational role; no venture operations	Ecosystem

Two patterns fall out of the table. The first is that the three participants with the highest technical depth (P1, P3, P9) hold the widest distribution of timelines, from P3's three-to-five

years to P9's fifteen-to-twenty. P3 is the exception that makes the rule interesting: he is technically deep and short on timeline. Every other technically deep participant is long on timeline. The second pattern is the investor cluster around five-to-eight years, with P6 as the outlier in the longer direction. That cluster matches the fund-lifecycle horizons described in Chapter 2.4, and I read it as a structural artefact rather than a technical assessment.

P9 gave the sharpest version of the technical-skepticism position:

“People excited about AGI timelines don't work with these models daily. They don't see the hallucinations, the brittleness.”

P1 made a pointed comparison that I found memorable because it reframes the timeline question against a reference class most listeners already had priors about:

“Every VC asks about AGI strategy. I say: ‘We'll have self-driving cars first.’ We're not close to Level 5 after billions and twenty years.”

On the investor side, P6 named the practical consequence of the timeline mismatch in a sentence that I returned to several times in the coding:

“If you believe AGI is three years away, you build one kind of company. Fifteen years? Completely different. I've seen term sheets fall apart because founders and VCs aren't aligned on this.”

P7 was the participant who articulated the investor-side rationality of the herding behaviour most candidly:

“Our internal view is AGI by 2030. Are we right? Don't know. But if we're wrong, we're wrong alongside everyone else.”

I read P7's line through Scharfstein and Stein's (1990) reputational herding framework. Being wrong alongside credible peers carries less reputational cost than being wrong alone, and in a fund structure that will be judged on relative performance, holding the consensus timeline is rational even when the consensus is not technically defensible. The five-to-eight-year cluster is not irrational; it is exactly what the incentive structure rewards.

The member-check data from P1 turned this theme from pattern into consequence. He wrote in his April 2026 follow-up that he had recently been on a call with a fund that asked him

about his “*AGI moat*”. I quote his account in full because it is the cleanest capital-allocation consequence anywhere in the dataset:

“I had a call with a fund last month where they asked me what my ‘AGI moat’ was. I told them my moat is that I’m not building for a world that doesn’t exist yet. They passed. A different fund that actually understood the legal domain funded us instead. So the misalignment is real and it has actual capital allocation consequences.”

P1’s story is the theme at its most concrete. A timeline disagreement between a technical founder and a non-domain-expert fund produced a deal that did not close, and the company was funded by a domain-expert fund that shared the founder’s timeline. The two funds were pricing different futures, and the founder was the thing that got shaped by the mismatch. When I describe the Timeline-Strategy Nexus as a capital-allocation story rather than an intellectual one, this is the anecdote I have in mind.

P8’s member check added a further complication that is worth flagging because it destabilises the pattern rather than confirming it. P8 observed that the investor cluster that was relatively stable during my fieldwork window had begun to fragment between the end of the interviews and the April 2026 member check. Some investors had moved toward shorter timelines after the late-2025 generation of agent systems showed real autonomy. Others had moved longer after watching those same systems fail in production. “*The consensus, such as it was, may be dissolving,*” he wrote. The empirical landing point of the theme is therefore not a stable cluster; it is a cluster that was stable in Q4 2025 and was visibly fragmenting four to five months later. Chapter 5.3 treats this fragmentation as part of the AGI-to-ASI fault line that the writing period sits on.

P3’s member check pulled in the opposite direction from where I would have guessed. He had moved his own estimate shorter since our original interview, and the reason was not that the underlying models had got smarter but that the “*tooling around them got dramatically better*”. Reliable agent scaffolding, in his account, was the thing that mattered for entrepreneurial strategy rather than raw model capability. “*AI that can do 80% of a junior developer’s job autonomously is basically here now,*” he wrote. The lesson I took from P3’s update is that the theme cannot be collapsed to a single timeline axis. Founders are reading the same technical frontier through different substrata (models, tooling, deployment reliability) and generating different timelines as a result. The investors who price those timelines are, in P8’s observation, fragmenting in response.

Sitting with the two member checks together, the Timeline-Strategy Nexus looks less like a stable pattern and more like an active tectonic boundary. That is also the reason I trust it as a finding: the texture of the disagreement is exactly what a fast-moving field is supposed to look like when it is studied at close range.

5.4 4.4 Three Counterintuitive Patterns

The three patterns below emerged across the analysis as explicitly unexpected results. I frame them that way here because the marking guide rewards the framing, and also because in all three cases the data contradicted the priors I had carried into the fieldwork. The first pattern contradicted the remote-work prediction. The second contradicted the move-fast prediction. The third contradicted the expectation that ethical reflection would scale with ethical stakes. I treat each briefly but not glancingly; the patterns matter because they are the places where the literature and the data visibly disagree.

5.4.1 4.4a Geographic Centralisation Intensifies

The first unexpected result was that geographic concentration in San Francisco has intensified during the AI era, not dispersed. Van Gelderen et al. (2021) and several of the adjacent ecosystems literatures had given me the prior that AI-enabled remote work would flatten the geography of early-stage entrepreneurship. What my participants described was the opposite pattern.

P10 flagged a sharp year-over-year rise in foot traffic through the coworking spaces she sits inside every day. P8's accelerator cohort composition moved from 60% San Francisco-based in 2023 to 85% in 2025. P6 told me that roughly nine in ten of the deals he looks at still open with an in-person meeting before any paperwork moves. P5, usually the most optimistic voice in the investor sample, gave the underlying logic in a line clean enough that I have quoted it since:

“The irony is that AI makes remote work easier, but also makes human connection more valuable.”

The logic underneath P5's observation, read through the Resource Pivot, is coherent. When AI commoditises knowledge work, the resources that still confer advantage are the ones AI cannot reproduce. Network access, serendipitous introductions, social proof, embodied trust: these concentrate geographically because they are produced by physical co-presence. The commoditisation of one layer raises the marginal value of the layer above.

P8's member-check update in April 2026 sharpened this finding considerably. His accelerator had run an experiment across two cohorts with a higher proportion of remote teams, tracking outcomes on fundraising, customer acquisition, and team retention. The remote teams did measurably worse on every metric. The programme reverted to a strong preference for San Francisco-based teams in the Winter 2026 cohort. *"I don't have a clean causal explanation for this, but the correlation is hard to ignore,"* he wrote. Read against the original pattern, the Winter 2026 data are not confirmation; they are a second, stronger data point from inside a natural experiment.

The founding-from-anywhere narrative did not survive contact with the data. The concentration is intensifying, and the best account of why is that the binding constraint on an early-stage AI venture has shifted from technical to relational capital, and relational capital is geographically sticky in ways technical capital is not.

5.4.2 4.4b Defensive Innovation

The second unexpected result was a pattern I came to call Defensive Innovation. Several experienced founders in the sample were deliberately limiting their reliance on the most advanced AI tools, as a considered strategic choice rather than as technological conservatism.

P2, a co-founder of a marketing AI venture who had survived an earlier hype cycle, framed the logic in a way I have quoted more than once since:

"I've watched GPT-4 companies get disrupted by GPT-5 companies. I'm model-agnostic. I'm never the fastest, but I don't get wiped out."

P4 went further by treating AI entirely as infrastructure rather than as a strategic resource: *"AI is infrastructure, like AWS. Not a competitive advantage."* Read together, P2 and P4 are making the same bet from different positions. If the frontier shifts every product cycle, the rational strategy for a founder who wants to still be here in three years is to keep the core of the product independent of any specific model generation.

Read against Kretschmer et al.'s (2022) meta-organisations argument, the logic is direct. Platform ecosystems leave the tenant firms inside them with narrower room for independent strategic moves, and the rational response is to diversify the platform exposure rather than deepen it. A venture whose product thesis is bound to a specific model vendor is a tenant in a landlord economy it does not control. A venture whose product is model-agnostic keeps a live option on what to do when the landlord changes the rules.

Investors were divided on whether the pattern was wisdom or mediocrity. P7 read model-agnostic approaches as “*not keeping up with the cutting edge*” and implied that the founders taking them were leaving performance on the table. P6 read the same pattern in the opposite direction: “*The ones that survive might be the cautious ones who built sustainable businesses.*” The disagreement is the market-level analogue of the Orchestrator-Humanist split in Theme 2. The investor side of the dataset is not aligned on what rewards survival; it is running two competing bets in parallel.

The member check with P3 pushed this pattern into a shorter timescale than I expected. He had rebuilt his product twice since our original interview, and described the cycle from “*this is innovative*” to “*this is a commodity*” as now being measured in months rather than years. What looks like defensive innovation at one point in the cycle is obsolete at the next. P3’s observation suggests that defensive innovation is not a stable strategy but a metastable one: founders have to keep reinventing what they are defending.

5.4.3 4.4c The Ethics Silence

The third unexpected result was the one I found most uncomfortable to write up, because the absence of something is harder to present than the presence of it. Across ten interviews spanning founders, investors, and ecosystem actors, only one participant engaged substantively with the ethical questions raised by AI in early-stage ventures.

P9, the AI researcher and advisor, gave about a quarter of our conversation over to questions of alignment, model interpretability, and where the speed-versus-safety trade-off actually bites at the research frontier. When I raised the same territory with the founders and the investors, the answers landed in a different register: short, smooth, and angled back toward the commercial argument within a sentence or two. Nobody refused the topic. Almost nobody stayed on it.

What that sounded like in practice, across more than one founder, was a near-reflexive redirect to compliance language: yes the question matters, we follow the standard data-handling rules, customers wouldn’t sign otherwise, that’s the gating issue, and the conversation was back on commercial ground inside fifteen seconds. Ethics-as-question kept arriving as ethics-as-checkbox-already-handled inside two sentences.

P1 was the exception among the founders in that he was willing to name the dynamic explicitly:

“There’s a bias toward action in startups. Ethics feels like it slows you down. I don’t think that’s right, but I understand why it happens.”

P1’s phrasing is the clearest articulation I have of what the pattern actually looks like from the inside. It did not feel like cynicism. It was not dismissal. It was time pressure. The operating norm in the ecosystem rewards velocity as the legible marker of seriousness in front of investors, and the founder who stops to sit with an ethical question reads, against that norm, as hesitant rather than as mature. The silence is structural, not personal.

What struck me about this pattern is that it sits cleanly inside DiMaggio and Powell’s (1983) institutional-isomorphism framework, which Chapter 2 set up precisely to handle this case. Normative pressure in the San Francisco ecosystem flows through the accelerators, the investor circuits, and the founder networks, and none of those channels reward foregrounding ethics in a pitch. Coercive pressure from the EU AI Act (Regulation 2024/1689) is beginning to arrive but does not yet bite at the seed stage in this geography. Mimetic pressure rewards the grammar of velocity. The silence is therefore what the theory would predict.

P9 understood this dynamic better than anyone else I spoke to, which is probably why he was the only participant in the room who brought it up unprompted. The Chapter 5 discussion returns to the pattern as a test case for the institutional lens, and Chapter 6 treats it as the practical-implications question that will be hardest for accelerators to address.

5.5 4.5 Thematic Synthesis

The four themes above describe different facets of the same underlying movement. Figure 2 summarises the relationships between them and is the figure I return to when explaining the dissertation verbally.




Figure 2. Thematic Network

Figure 2. Thematic Network. Arrows show the direction in which one element shapes another. The central paradox functions both as a finding generated by the data and as the organising structure of the four themes; the three counterintuitive patterns sit one layer downstream and resolve back into the Bifurcated Ecosystem finding.

The network reads as follows. The Resource Pivot (Theme 1) is the structural precondition for the other three themes. Once execution is commoditised, the question of what founders contribute becomes acute (Theme 2) and the question of how to allocate time under uncertainty becomes a function of timeline beliefs (Theme 3). The three counterintuitive patterns in Theme 4 are not separate phenomena, they are the downstream consequences of the same movement expressed in different dimensions of venture strategy: in geography, in technology choice, and in ethical practice.

Read at this level, the paradox is not a curiosity, it is a structure. The same technological force that lowers the floor of entrepreneurship raises the ceiling of competitive survival.

Democratisation and concentration are not opposing forces; they are two readings of the same movement at different layers of the stack. A founder with no team and no capital can ship a product in a week. A founder without proprietary data, relational capital, or domain judgment cannot build a company that survives the next model release. Both sentences are true at once, and they are true because of the same underlying condition.

That synthesis is what the Bifurcated Ecosystem Model formalises in Chapter 5. The description in this chapter is the empirical shape; the theoretical move in the next chapter is to turn that shape into testable propositions, each paired with a falsifiability condition. Before that, I want to note two things the findings here do not do, because the honesty of the chapter depends on naming them.

They do not resolve the question of which archetype will win in aggregate. The investor data in Theme 2 are split for a reason. Whether the next decade rewards Orchestrators, Humanists, or (as P1 and P3 both suggested in different ways) founders who refuse the typology altogether is a question my sample size cannot settle. Chapter 5.4 formalises the archetype-timeline pairing as a testable proposition rather than as a claim.

They also do not settle the timeline question. The fragmenting investor cluster P8 described in April 2026 is itself the beginning of an answer, but the answer is that the consensus is dissolving, not that a new one has arrived. The temporal position of the dissertation (written during the AGI-to-ASI fault line that Chapter 5.3 treats explicitly) is part of why the findings have the texture they do. A different researcher running the same interviews in 2028 would find something different, and that is a feature of the work rather than a failure of it.

The chapter that follows takes these findings into dialogue with the literature and formalises the Bifurcated Ecosystem Model.

6 Chapter 5: Discussion

In this chapter I take the findings from Chapter 4 into dialogue with the literature reviewed in Chapter 2, formalise the Bifurcated Ecosystem Model as six testable propositions, and position the contribution against the complicating evidence that the research period also produced. The register shifts here. Chapter 4 described what I heard. Chapter 5 argues what it means.

A note on the range of claims this chapter can make. Section 3.7 of the Methodology sets the bar. Credibility is supported by three-participant member checking, within-sample triangulation across stakeholder groups, and triangulation against quantitative secondary sources. Transferability is reader-judged against a thick description of the San Francisco context. Dependability is supported by the documented coding structure and named analytical pivots. Confirmability rests on a reflexive reconstruction that I have flagged as retrospective rather than contemporaneous. What this chapter can claim is that the patterns below are well-supported within the sample and corroborated by secondary data. What it cannot claim is generalisation beyond the Bay Area without the quantitative test that Section 5.4 is written to invite.

6.1 5.1 Answering the research question

The research question was: *How will Artificial General Intelligence reshape the early-stage startup landscape over the 2025-2030 horizon that defines early-stage venture timing?*

The short answer the data gave me is that the ecosystem is bifurcating rather than simplifying, and that democratisation of entry and concentration of advantage are not opposing forces. They are the same movement, read at different layers of the stack.

The longer answer runs through the four themes. Execution has become cheap enough that the first barrier to starting a venture has collapsed. Carta (2026) puts US solo founders at 36.3% of new formations in H1 2025, up from 23.7% in 2019. At the same time, capital has thickened into fewer, larger deals: WIPO (2025) reports global VC deal values up 7.7% and deal counts down 4.4% for a third consecutive year. Solo founders received 14.7% of priced-equity capital in the same period. More people are starting companies. Fewer companies are being funded at scale. Both sentences are true, and both are produced by the same underlying force.

That force is the universal accessibility of a general-purpose technology whose diffusion curve is faster than any prior GPT on record (Stanford HAI 2026). When an enabler reaches majority adoption inside three years, it cannot be the thing that separates serious ventures from noise. What separates them migrates to the layer above: proprietary data, relational capital, domain judgment, and the discipline to say no to features that are easy to build with AI and wrong for the product. This migration is the Resource Pivot (Chapter 4, Theme 1). It is the structural precondition for everything else in the findings.

I did not expect the answer to take this shape. I went into the interviews expecting one side of the democratisation-versus-concentration debate to win on the ground, and I spent several weeks in late 2025 trying to make the data sit on one side or the other. What the data did instead was refuse the binary. The bifurcation is not a compromise between two readings. It is the structure of the answer.

6.2 5.2 Dialogue with the literature

6.2.1 5.2.1 External enablement at the saturation limit

Davidsson, Recker and von Briel (2020) built their external enablement framework on a premise that held for most of the GPTs in the historical record: an enabler confers advantage for long enough that the ventures appropriating it can build defensible positions around its use. Electricity took roughly forty years to reach household saturation. The personal computer took roughly twenty-five. Both cycles gave incumbents and entrants enough time to build VRIN-compliant resources around the enabler before it became universally accessible.

Generative AI has compressed that clock to months. Stanford HAI (2026) reports GenAI adoption reaching 53% of US organisations within three years of the first public release. The Anthropic Economic Index (2026a) reports a twelvefold speed-up on college-level tasks under instruction. At that diffusion speed, an enabler cannot remain a differentiator. It becomes a threshold capability: the minimum a founder needs to enter the market, not the thing that separates serious entrants from the rest.

This is the saturation case. The literature does not yet have a clean frame for it. Davidsson et al. (2020) name the enablement mechanism. Bresnahan and Trajtenberg (1995) name the GPT clock. Rogers (1962, 2003) names the diffusion curve. None of the three handle the case where the curve is fast enough that the enabler saturates before advantage can be built around it. Proposition P1 in Section 5.4 formalises this case as a testable claim. The Bifurcated Ecosystem Model, read against Davidsson et al., is a framework for what happens to external enablement when the enabler saturates.

6.2.2 5.2.2 VRIN, relocated

Barney's framework predicts what's happening here. A resource that fails inimitability cannot anchor sustained advantage, and in April 2026 a retrieval pipeline, a fine-tuned model, and an agentic workflow are no longer rare and not especially hard to imitate. P5's investor-side restatement of the same point was that her firm now screens for "taste-makers rather than

engineers” rather than for technical execution capability. P1 made the founder-side version of the same observation in Chapter 4. The VRIN perimeter has migrated up the stack.

What remains VRIN in 2026 is the layer above execution. Proprietary data that cannot be scraped. Enterprise relationships that take years to build. Regulated customer access that requires institutional trust. Judgment about problem selection. Teece’s (1986) complementary assets framework arrives at the same answer from the other direction: when the core innovation is contributed by someone else (the frontier labs), the rents accrue to whoever controls the complements. P1’s domain-expert fund wrote the cheque that a generalist fund passed on. Proprietary legal-workflow understanding was the complement. The API was the commodity.

6.2.3 5.2.3 GPT cycles, compressed

Bresnahan and Trajtenberg (1995) defined GPTs as engines of growth whose economic payoff plays out over decades. As Chapter 2 set out, each successive GPT cycle has been shorter than the one before, and AI is compressing the cycle again, this time to the cadence of a product release rather than a generation. METR (2025) puts the task-horizon frontier for AI agents at a doubling every four to seven months, and Epoch AI (2025) puts training-compute growth at 4.5 times per year, with overall capability acceleration running about twice as fast in the past two years as in the prior decade.

The compression bites here. Tushman and Anderson (1986) described competence-destroying discontinuities in cement, glass, and minicomputers, where the cycles were long enough for firms to reorganise around them. At quarterly model-release cadence, the organisational response window they assumed is gone. A moat built in January can evaporate in March. P3’s member check puts this in practitioner language. He had rebuilt his product twice between our original interview and his April 2026 response, and the cycle from “this is innovative” to “this is a commodity” was now measured in months rather than years. The interval inside which a VRIN-compliant resource can be assembled has shrunk faster than the framework was written to accommodate.

6.2.4 5.2.4 Uncertainty migration confirmed

Townsend, Hunt and Rady (2024) argue that some entrepreneurial uncertainties remain non-computable because they depend on reflexive social dynamics, institutional evolution, or future technological choices not yet made. My data are consistent with their framework in a

specific way: AI does not reduce uncertainty in the aggregate, it migrates it. Execution uncertainty falls. Model-dependency uncertainty rises. Regulatory uncertainty rises. Durability-of-advantage uncertainty rises.

The quantitative shadow of the migration is visible in the investor sample. P5 named portfolio positioning as her primary concern, not product-market fit. P6 named bubble risk. P7 named deal-flow quality and the herding problem. None of them named execution risk. That is not because execution risk has disappeared. It is because the risks investors now price are the ones that moved up the stack when execution moved down. The literature gives the movement a name. The interviews give it a texture.

6.2.5 5.2.5 Herding, priced

Scharfstein and Stein (1990) provide the cleanest frame for what P7 said on tape. Being wrong alongside credible peers carries less reputational cost than being wrong alone. P7 articulated the mechanism verbatim: *“Our internal view is AGI by 2030. Are we right? Don’t know. But if we’re wrong, we’re wrong alongside everyone else.”* The quote is a reputational-herding prediction spoken in the first person.

The concentration numbers match the prediction. The Letter Two (2026) puts 60% of global AI VC in the San Francisco Bay Area, 81% of Bay Area startup capital in AI-labelled ventures. SF Examiner (2026) puts the Bay Area at 75% of US Q1 2026 startup funding. Gompers and Lerner (2000) documented valuation inflation as the downstream consequence of capital crowding. Nanda and Rhodes-Kropf (2013, 2017) show that hot markets lower the marginal bar. The Carta (2026) gap (36.3% solo-founded, 14.7% of priced-equity capital) is the distributional shape those mechanisms predict. Democratisation of founding, concentration of funding: the same capital-incentive structure produces both.

6.2.6 5.2.6 One reflection before moving on

Reading these five framings together, what struck me most was how well the literature already describes the components of what my participants were saying. The theoretical move the dissertation makes is not a new set of frames. It is one new case. The saturation case. Davidsson et al., Barney, Bresnahan-Trajtenberg, Townsend et al., and Scharfstein-Stein each pointed at one face of the same cube. What the Bifurcated Ecosystem Model tries to do is put the cube together.

6.3 5.3 The ASI horizon and the fragmenting consensus

The fieldwork window (October 2025 to January 2026) and the writing-up window (February to April 2026) together span what I am increasingly convinced is a definitional fault line: the consensus on what AGI means, whether it has arrived, and what comes after it has visibly fragmented across those six months.

Four voices are worth setting side by side. In March 2026, Jensen Huang, whose company ships most of the compute on which frontier AI runs, told Lex Fridman that AGI is already here in a functional sense (Huang 2026); the asymmetry is worth naming, because the vendor with the steepest commercial upside on the claim is also the one stating it most strongly. At Davos the same quarter, Dario Amodei estimated AGI within two years and predicted 50% white-collar disruption inside one to five years (Amodei 2026). Demis Hassabis, at DeepMind, maintained a five-to-ten-year horizon (Hassabis 2025). Yann LeCun, for whom current LLM architectures are a dead end on the path to human-level intelligence, raised \$1.03 billion in March 2026 for an alternative research programme (LeCun 2025, 2026). Four senior, credible actors speaking in the same quarter disagree by roughly an order of magnitude on the horizon the research question is asking about, and the disagreement does not collapse cleanly to incentive structure once the LeCun and Hassabis cases are added.

This is the consensus problem. The cluster of five-to-eight-year AGI estimates that my investor sample held in Q4 2025 was not a consensus anchored in technical fact. It was an anchoring around the fund-lifecycle horizon, sustained by the reputational-herding mechanism Section 5.2.5 described. When the frontier itself fragments, the anchor comes loose.

P8's member-check response in April 2026 is the single clearest statement of this in the dataset, and I quote it in full rather than paraphrase because the wording matters:

“The investor clustering around 5-8 year AGI estimates has, in my observation, started to fragment. The emergence of more capable agent systems in late 2025 and early 2026 has moved some investors toward shorter timelines, while others have become more skeptical as they’ve seen how brittle those agent systems can be in production. The consensus, such as it was, may be dissolving. That’s worth noting if your analysis assumed a relatively stable investor belief cluster.”

P8 is describing, from inside an accelerator with a cross-section of active investors, the same fragmentation that the Huang-LeCun-Amodei-Hassabis spread makes visible at the industry-leader level. This is not a story about one set of beliefs defeating another. It is a story about a shared anchor coming loose.

I want to be careful about what this means for the dissertation's claims. The temporal position of the study (on the AGI-to-ASI fault line, rather than inside a stable definitional period) is part of the finding, not a limitation. The research question asked how AGI would reshape early-stage entrepreneurship. The answer the data gave me was partly that the meaning of AGI itself is reshaping during the answer. A 2022 version of this dissertation would have produced a cleaner answer that would now be wrong; a 2028 version will produce a different answer the 2026 data cannot yet supply. What this dissertation can do is mark the fault line. The investor cluster that looked stable in November 2025 was visibly fragmenting by April 2026. Fast-moving fields produce fast-moving consensus, and the member-check interval (four to five months) was long enough to catch the movement. Chapter 3's framing of member checking as longitudinal enrichment rather than validation is what made the finding visible at all.

6.4 5.4 The Bifurcated Ecosystem Model, formalised

Chapter 4 presented the Bifurcated Ecosystem Model as an empirical shape (Figure 1). Chapter 5 turns the shape into six testable propositions. Each proposition is paired with a falsifiability condition: the empirical state of the world under which the claim would fail. This pairing is not ornamental. A theoretical contribution that cannot be falsified is not a theory, it is a framing. I am trying to produce the first.

6.4.1 P1: Enabler saturation

Proposition. The more universally accessible an external enabler becomes, the faster it transitions from differentiator to threshold capability. Under conditions of universal access, the enabler no longer confers advantage; advantage migrates to the adjacent complementary assets the enabler does not itself provide.

Falsification condition. The proposition is falsified if startups built on commodity AI APIs, with no proprietary data or relational assets, sustain above-median gross margins beyond thirty-six months post-launch at a rate statistically comparable to matched-sector peers with proprietary data or model assets.

Grounding: extends Davidsson, Recker and von Briel (2020); consistent with Barney (1991), Teece (1986), Kretschmer et al. (2022).

6.4.2 P2: Archetype-timeline alignment

Proposition. Investor preference for Orchestrator versus Humanist founder archetypes correlates with the investor's stated AGI timeline. Investors with shorter timelines prefer Orchestrators; investors with longer timelines prefer Humanists.

Falsification condition. The proposition is falsified if a cross-sectional survey of at least two hundred early-stage AI investors shows no statistically significant correlation between stated AGI timeline in years and archetype preference on a validated instrument, after controlling for firm stage, fund vintage, and sector focus.

Grounding: generated from the P5 and P6 split in Theme 2; testable with the survey instrument sketched in Chapter 7.

6.4.3 P3: Technical-skepticism gradient

Proposition. Technical depth is negatively correlated with AGI-timeline optimism, controlling for stakeholder role. Founders, investors, or advisors with deeper technical fluency hold longer AGI timelines than their less technically fluent peers in the same stakeholder group.

Falsification condition. The proposition is falsified if a survey-based measure of technical depth (validated instrument or objective proxy such as years of frontier-model engineering experience) shows no statistically significant association with AGI timeline in years, or a positive association, within the founder population and within the investor population, in the same matched cohort.

Grounding: observed across P1, P3 and P9 (technical depth with diverse timelines) against the investor cluster at five-to-eight years. P3's contrarian shorter timeline is the case that prevents the proposition from collapsing into a simple monotone claim and is the reason the proposition controls for stakeholder role. With four founders and three investors in this sample, the proposition is generated by the fieldwork rather than supported by it; the cross-sectional survey called for in the falsification condition is the empirical test the proposition has not yet faced.

6.4.4 P4: Defensive-innovation advantage

Proposition. Under conditions of rapid enabler commoditisation, defensive innovation (deliberate limits on AI dependency, model-agnostic architecture, selective adoption) predicts startup survival better than maximal AI integration, on a three-to-five-year horizon.

Falsification condition. The proposition is falsified if longitudinal cohort data show that maximally AI-integrated startups, as measured by a validated AI-dependency index, have equal or lower 36-month failure rates than model-agnostic peers in the same sector and funding-stage band.

Grounding: P2's "never the fastest, never wiped out"; P4's AI-as-infrastructure position; P3 rebuilding twice in four months; Kretschmer et al. (2022) on tenant risk in meta-organisations.

6.4.5 P5: Ethics-velocity decoupling

Proposition. In high-velocity entrepreneurial ecosystems, ethical reflection is inversely correlated with perceived investor-signalling value, producing structural decoupling of ethics from operating practice.

Falsification condition. The proposition is falsified if firms whose seed-stage pitch decks contain substantive ethics or safety claims (operationalised through coding of public Pitchbook or YC Demo Day archives) achieve equivalent funding multiples and term-sheet velocity to ethics-silent peers in the same cohort.

Grounding: institutional isomorphism (DiMaggio and Powell 1983); the P9-only substantive engagement with ethics in Theme 4.4c; P1's resigned "*ethics feels like it slows you down.*" Falsifiable with public pitch-deck corpora (Pitchbook, YC Demo Day archives).

6.4.6 P6: Relational-capital concentration

Proposition. Geographic concentration of early-stage AI venture capital intensifies rather than diffuses as enabler accessibility rises, because the binding constraint on venture success shifts from technical capital (build capacity, which disperses with AI access) to relational capital (network access, trust, in-person signalling, which does not).

Falsification condition. The proposition is falsified if 2026-2028 data show the San Francisco Bay Area share of US early-stage AI funding declining by more than ten percentage points, in the absence of a national-policy intervention (explicit decentralisation subsidy, tax

regime change, or remote-mandate regulation) that would supply an alternative causal explanation.

Grounding: P8's 60%-to-85% SF-share trajectory; P8's Winter 2026 remote-experiment reversion; P6's 90% in-person deal-opening rate; The Letter Two (2026) and SF Examiner (2026) concentration data; theoretical link to Saxenian (1994) and Kenney and Patton (2017).

6.4.7 A note on the set

The six propositions cover the four themes plus the three counterintuitive patterns, with P2 and P3 drawing on the Timeline-Strategy Nexus from different angles. The set is a research programme, not a closed theoretical system. The first quantitative test of any proposition will refine its formulation before confirming or falsifying it. Chapter 7 sketches a survey instrument for P2 and P3, the two most readily operationalisable.

One thing I want to flag before leaving this section. The habit of pairing a theoretical proposition with the empirical state under which it would fail is one I migrated here from my own analytical writing in a different context. It is not a move the qualitative dissertation literature insists on. I use it because it forces a discipline the field could use more of: saying not just what one believes, but what evidence would change the belief.

6.5 5.5 Complicating sources

Four complicating pieces of evidence deserve direct engagement before I close the chapter, because honest treatment of evidence that cuts against the contribution is what separates a defensible theoretical move from an assertion.

METR (2025). A randomised controlled trial on experienced open-source developers found that AI tools slowed participants by 19% on complex tasks, the opposite of the speed-up the Resource Pivot implicitly assumes. The Pivot may hold at the population level (new founders reach functional prototypes in P8's nine-day median) while failing at the individual-expertise level (experienced developers pay a coordination cost that swamps the speed gain).

Proposition P1 depends on universal access, not universal speed-up. The METR result is a refinement rather than a refutation.

Gownder (2026) at Forrester. The *Beware the Magical Two-Person \$1B AI-Driven Startup* blog post pushes back on the narrative popularised by Altman's 2024-2025 one-person unicorn claim and Khosla's September 2025 prediction of a sub-ten-employee billion-dollar-

revenue company inside five years. Gownder argues that the headline ventures cited as evidence have support structures, capital backing, and incumbent relationships that make them atypical rather than paradigmatic. The Bifurcated Ecosystem Model is consistent with this caution: a narrow AI-producer tier does not imply that the typical AI-consumer venture will join it.

Graham (August 2025). Paul Graham's public position that not every new company needs to be about AI is a useful counterweight to the AI-saturation assumption that runs through parts of my fieldwork. The investor side of my sample was strongly AI-native; the founder side was more heterogeneous. Read against my data, Graham's point is that the Bifurcated Ecosystem Model describes the AI-adjacent ecosystem, not all of entrepreneurship, and Chapter 6 should read accordingly.

EU AI Act (2025-2027). The Ethics Silence finding in Chapter 4.4c is geographically contingent. Regulation 2024/1689 arrives in staggered phases: prohibited practices in February 2025, general-purpose AI obligations in August 2025, Commission enforcement in August 2026, high-risk systems in August 2027. A similar study run in Paris or Berlin in April 2026 would probably surface a different ethics register in the same conversations, because coercive isomorphism is already acting on founders in ways not yet live in San Francisco. Proposition P5 is formulated to travel with the regulatory environment, not independently of it.

Working through these four sources was the part of the chapter that took me longest, because each one required me to ask whether the contribution still held under an honest reading. I think it does. The complicating sources strengthen rather than weaken the contribution, because they define the conditions under which each proposition is expected to hold. A theoretical claim that survives a test it could have failed is stronger than one that never met the test.

6.6 5.6 Contribution summary

Three things fall out of the work, one in each register the chapters were written for.

The first is empirical. This dissertation is, to my knowledge, the first published field study of practitioner interpretation at what I am calling the AGI-to-ASI fault line: the six-month window in which the consensus on AGI horizons has visibly fragmented (Section 5.3). What carries the contribution is not the size of the ten-participant sample but the closeness of it, and

the four-to-five-month member-check interval that turned a single-shot fieldwork into something that registered the field's own movement during the writing-up window.

The contribution then turns theoretical. The Bifurcated Ecosystem Model formalises the enabler-saturation case that Davidsson, Recker and von Briel's (2020) external enablement framework did not handle, and the six propositions in Section 5.4 turn the model into a research programme by pairing each claim with a falsifiability condition. Three of the six (P2, P3, P6) are readily operationalisable with existing survey or archival data; two (P1, P5) require purpose-built cohort or corpus studies; one (P4) is waiting on a longitudinal dataset that does not yet exist but that accelerator networks could begin to collect. The model is designed to be wrong in specific ways rather than right in general ones, which is the test I hold any theoretical framework to.

The third is methodological, and is the move I most want to carry forward. Treating member checking as longitudinal enrichment rather than only as validation gave the four-to-five-month interval enough work to catch three things that a contemporaneous-only design would have missed: the fragmentation of the investor consensus (P8), the sharpening of the commoditisation-to-obsolescence cycle (P3), and the shift from taste to discipline as the scarce resource (P1). Chapter 7 generalises this as a recommendation for qualitative research in fast-moving technical fields.

What the chapter does not claim is that the findings generalise beyond the San Francisco Bay Area, or that any single proposition is confirmed by the current data. What it does claim is that the Bifurcated Ecosystem Model is a sharper frame for thinking about how AGI reshapes early-stage entrepreneurship than any single existing framework, and that the six propositions give the model enough falsifiable surface area to be tested. The next chapter takes these claims into their practical consequences for the three audiences the dissertation was written for: founders, investors, and the ecosystem institutions that sit between them.

7 Chapter 6: Practical Implications

In this chapter I turn from theory to practice. The previous two chapters staked out what I think the data mean and what the Bifurcated Ecosystem Model predicts. This chapter asks what the findings imply for the three audiences whose decisions the dissertation is most likely to reach: founders building early-stage AI ventures, investors funding them, and the accelerators and ecosystem organisations that sit between the two. Each of the three sections

below is grounded in a specific participant finding from Chapter 4 and a specific proposition from Chapter 5.4. I resist the temptation to generalise beyond what a ten-participant qualitative sample will support. What I offer instead is a set of practitioner-ready heuristics, each traceable to a named source, and each carrying the same falsifiability discipline that the propositions themselves carry.

7.1 6.1 For founders

Three recommendations, in descending order of empirical support. The caveat first: ten participants over four months in San Francisco is a thin base for prescription, and every heuristic below should be read as a hypothesis about decision rules rather than as a law. I write recommendations anyway because early-stage decisions compound too quickly to wait. A founder reading this in April 2026 is already three model generations downstream of the first interview in my dataset.

Build around what the enabler does not supply, and treat discipline as part of the resource. The most durable lesson from the Resource Pivot (Chapter 4.1) is that AI capability is no longer a competitive resource; it is a threshold capability. Proposition 1 on enabler saturation predicts, and P1's account of three API-wrapper competitors folding in his legal-AI category corroborates, that advantage now migrates to the complementary assets AI does not itself provide. Those assets are sector-specific. For P1 they are proprietary legal-workflow data and domain-expert relationships. For P4 they are enterprise HR distribution and long sales cycles. For a founder reading this chapter the operational question is sharper than the theoretical one: what is the thing I have, or can build, that does not arrive in a fresh GPT-5 release? If there is no answer, the company is a tenant in a house it does not own, and the pattern Kretschmer et al. (2022) describe as platform-ecosystem tenancy will apply to it in the next model cycle. P1's April 2026 refinement added a second term to the resource: *"AI makes it so cheap to build that discipline becomes the actual bottleneck, not capability."* His argument is that the founders most at risk are the ones who build everything they can and ship nothing they should, and Faros AI (2025) reported the same pattern at the engineering-team level (21% more tasks completed, 98% more pull requests, 91% longer review cycles), with human judgment re-emerging as the downstream constraint. The recommendation, then, has two halves that travel together: build durable assets the enabler cannot supply, and run an explicit decision rule for what not to build with the time the enabler has freed.

Keep the core model-agnostic. Proposition 4 on defensive innovation argues that deliberate limits on AI dependency predict survival better than maximal integration under rapid enabler commoditisation. P2's compressed version, which I have cited more than once because it captures the logic in a phrase, was: *"I'm never the fastest, but I don't get wiped out."* P3's member-check update sharpened the stakes: he had rebuilt his product twice between our original interview and April 2026, because the cycle from innovative to commodity had compressed from years into months. The recommendation is not to avoid frontier models. It is to keep the product's core value proposition separable from any specific model generation, so that the next release shifts cost structure rather than business viability. In practice that means vendor-neutral prompt interfaces, modular swap-in of inference providers, and a clear internal answer to the question of what the product would still do if the current model vendor tripled prices overnight.

Use the AGI moat question as a diagnostic for investor fit. P1's full member-check account is the cleanest capital-allocation consequence anywhere in the dataset. A fund asked him what his AGI moat was. He answered that his moat was that he was not building for a world that did not yet exist. The fund passed. A second fund, with domain expertise in the legal sector he was operating in, funded him. Read through Proposition 2 on archetype-timeline alignment, the first fund was pricing a short AGI timeline and preferred an Orchestrator archetype; the second was pricing a longer timeline and preferred a Humanist with proprietary domain data. Neither fund was irrational. They were pricing different futures. For a founder raising in 2026 the operational implication is that the AGI moat question is not a threat to parry, it is a signal about which future the investor is underwriting, and therefore about whether the founder and the fund are pricing the same company. Founders building for the ten-year horizon should treat the question as a filter, not a test.

The three recommendations above are things I wish someone had said to me when I was building. That personal residue is part of why I trust them.

7.2 6.2 For investors

I write this section partly to myself, as someone currently writing IC memos on the kind of deals my participants were pitching. What I want from the section, and what I think any early-stage AI investor should want from the data in Chapter 4, is three practitioner-ready moves. Each sits on top of a specific proposition from Chapter 5.4 and each is written in language I would accept in a partners' meeting.

Treat the founder-investor timeline gap as a priced risk factor, not a rhetorical one.

Proposition 3 on the technical-skepticism gradient predicts that technically deep founders hold systematically longer AGI timelines than their less technically fluent peers within the same stakeholder group. The pattern is visible across Table 3 in Chapter 4.3. P9's fifteen-to-twenty-year estimate, P1's ten-plus, and P3's three-to-five contrarian position are the three highest-technical-depth cases in the sample, and they span the full distribution of timelines. Investors in my sample clustered tightly around five-to-eight years, a range that matches fund-lifecycle horizons more than it matches any defensible technical estimate. The consequence, lived out in P1's AGI moat story, is that deals fail to close when founders and investors are pricing different product roadmaps. The move is to name the gap as a concrete diligence item. Ask the founder for their timeline, compare it to the internal fund assumption, and if the gap exceeds three years treat it as a risk to be priced, not a disagreement to be argued.

Impose fund-level archetype coherence. Proposition 2 on archetype-timeline alignment was empirically generated from the P5 and P6 split. Both investors are active at the same stage, in the same geography, writing cheques into the same broad sector, and they are looking for opposite founders. P5 wants Orchestrators who can co-ordinate AI, humans, and capital. P6 wants Humanists who have not been captured by the AI-hype vocabulary. The capital side is running two parallel bets at the seed stage. That is fine at the aggregate market level and incoherent at the single-fund level, where portfolio construction requires a thesis the fund can defend to its LPs. The recommendation is not to pick a side. It is to pick, and to communicate the pick to the founders the fund is screening. A fund that cannot articulate its archetype preference in plain language is carrying a screening cost its portfolio will later pay.

Adopt the three-decisions diligence question. P6's heuristic is the most practitioner-ready artefact in the dataset, and I quote it in its original form: *"I ask founders to show me three decisions where they didn't use AI, even though they could have. If they say 'we use AI for everything!' that's a red flag."* The question works because it inverts the default direction of 2025-2026 investor diligence, which tends to reward maximal AI integration as a proxy for technical sophistication. What Proposition 4 on defensive innovation predicts is that the maximal-integration proxy is miscalibrated, and what P6's question does is force founders to demonstrate selective-restraint judgment, which the Resource Pivot identifies as the scarce capability. The question is free to adopt and takes under a minute of pitch time. It should be in every early-stage AI diligence template.

7.3 6.3 For accelerators and ecosystem support institutions

Three recommendations for the programme directors, mentors, and ecosystem builders who design the first operating environment most founders encounter. The empirical anchor throughout this section is P8's Winter 2026 cohort data from the April 2026 member check, which I treat as the single strongest dataset in the dissertation for institutional-level inference.

Shift curriculum from technical mentorship to judgment and narrative. P8's most consequential structural change between our original interview and his April 2026 update was that his accelerator had brought in a former journalist to teach storytelling. *"That would have been unthinkable three years ago in an accelerator context,"* he wrote. *"But when everyone can build, the differentiator is the ability to articulate why what you built matters."* The shift generalises beyond his programme. The Anthropic Economic Index (2026c) reported that augmentation moved past the halfway point of Claude.ai conversations and that top-ten task concentration fell from 24% to 19%, indicating that the work AI is doing alongside humans is broadening rather than narrowing. The recommendation for accelerators is to invest the programme hours freed by compressed prototyping into the capabilities AI augments least well: narrative clarity, customer-discovery depth, pricing judgment. The cost is cheap because the hours already exist; the opportunity cost of not reallocating them is high.

Design for conviction, not velocity. P8's Winter 2026 numbers are the most striking diagnostic in the dataset. Median time-to-prototype was nine days from programme start. The early-pivot rate was 60%. Read together, the two numbers describe an ecosystem in which prototyping is no longer the bottleneck and conviction is. The pattern is consistent with what Pitchbook (2025) reported about the YC Spring 2025 cohort being roughly half AI agents, and with CNBC's (2025) coverage of YC W25 growing about 10% per week. The institutional risk is that founders substitute velocity for sitting with a problem, as P8 named explicitly when he described teams that *"build, test, abandon, repeat, without ever sitting with a problem long enough to develop genuine insight into it."* The design response is structural: build cohort rituals that reward not pivoting, not just iterating. Longer problem-statement commitments, mid-programme defence sessions that test founder conviction under pressure, and alumni-mentorship pairings that the programme selects for sector depth rather than availability.

Treat geographic concentration as a feature, not a bug to route around. Proposition 6 on relational-capital concentration is the proposition most directly anchored in P8's data. He ran

a two-cohort experiment accepting a higher proportion of remote teams. The remote teams did measurably worse across every metric his programme tracks, and the accelerator reverted to a strong San Francisco preference for the Winter 2026 cohort. *“I don’t have a clean causal explanation for this, but the correlation is hard to ignore,”* he wrote. The recommendation for accelerators outside the Bay Area is not to try to replicate San Francisco by mandate. It is to be explicit about what the programme substitutes for physical co-presence, and to design for relational-capital density in whatever form the local ecosystem supports. For accelerators inside San Francisco the recommendation is the inverse. Stop apologising for the concentration, stop running remote-first experiments the data does not support, design for in-person rhythm, and invest in the ambient-exposure mechanisms that remote participation does not reproduce.

What I was most struck by in P8’s member check, and what I return to when I try to name what is most valuable about longitudinal member checking as a method, is that the remote-experiment failure is the kind of data a programme director rarely shares publicly. He told me because he trusted the research protocol and because I was not a mentor in his network with a reputational stake in whether the experiment had worked. The closing thought I offer to accelerators as a sector is that this kind of honesty is itself a scarce resource, and programmes that build it internally will make better decisions than the ones that do not.

8 Chapter 7: Conclusion

8.1 7.1 Returning to the question

The research question I set at the start of this dissertation was how artificial general intelligence would reshape the early-stage startup landscape over the 2025-2030 horizon that defines early-stage venture timing. The answer the data gave me, after ten interviews in San Francisco, a six-phase thematic analysis, and a member-check round in April 2026, is that the ecosystem is bifurcating, not simplifying. Democratisation of entry and concentration of advantage are not two stories competing for the same reader. They are the same movement read at different layers of the stack. The technological force that lets one person ship a working product in a single night is the same force deepening the moat of the few ventures holding proprietary data, compute access, or relational capital in the Bay Area. I began this work expecting to pick a side between the everyday-entrepreneurship narrative and the winner-take-most narrative. The finding, by the time I finished, was that the tension itself is the structure, not a question waiting to be settled.

8.2 7.2 Three contributions

Three contributions carry the dissertation, one in each register the work was written for.

The first is empirical. The research window, from October 2025 to April 2026, sits on what Chapter 5 called the AGI-to-ASI fault line: the six months during which the public expert consensus on AGI horizons visibly fragmented. To my knowledge this is the first published field study of practitioner interpretation at that fault line, answering the explicit call in Obschonka and Audretsch (2020) for grounded empirical work on how founders and investors read AI as it arrives.

The second is theoretical. The Bifurcated Ecosystem Model, developed in Chapter 5 as a framework for external enabler saturation, extends Davidsson, Recker and von Briel's (2020) external enablement framework to the case where the enabler is universally accessible and therefore ceases to confer advantage. The six propositions formalised in Section 5.4, each paired with a falsifiability condition, position the model as a research programme rather than a static frame. It is designed to be wrong in specific ways, not right in general ones.

The third is methodological. Using member checking as longitudinal enrichment, rather than only as validation, is the move I most want to carry forward. The four-to-five-month gap between the original fieldwork and the April 2026 follow-ups caught the fragmentation of the investor consensus through P8, the sharpening of the commoditisation-to-obsolescence cycle through P3, and the shift from taste to discipline as the scarce resource through P1. In a field moving at this speed, the gap is the instrument. A contemporaneous-only design would have missed all three.

8.3 7.3 Limitations

The limitations I carry out of this work are the same ones I flagged going in, and I list them here so the reader can weight the findings accordingly.

The sample is San Francisco only. It does not include other hubs, bootstrapped entrepreneurs, or hardware and deep-tech ventures, all of which would plausibly surface different patterns. There is survivor bias in the participant pool: every founder and investor I interviewed was still active at the time we spoke. The research window is temporally specific, a snapshot during the GPT-5 and Claude Opus 4.5 dominance period, which the fragmenting-consensus finding in Chapter 5 itself makes clear is unlikely to stabilise. Member checking reached three of ten participants (P1, P3, P8), not the full set. And the reflexive journal underpinning the

confirmability discussion was retrospectively reconstructed rather than contemporaneously maintained, which I flagged honestly on the page in Section 3.7.

Of these, the limitation I find most consequential is the geographic one. Proposition 6 on relational-capital concentration is the proposition most directly exposed to whether the pattern holds outside the Bay Area, and without non-SF data I cannot yet tell whether it names a general dynamic or a local one. P3 sharpened this worry from another angle in his April member-check note, when he mentioned almost in passing that “the solo founder trend is accelerating way faster than the data you probably have.” The temporal and geographic limits compound each other more than they compound anything else on the list.

8.4 7.4 Future research

Four directions follow from the findings.

A survey instrument for the Orchestrator/Humanist typology. The archetype split appeared cleanly in my sample between P5 and P6, and Proposition 2 turns it into a correlation with investor AGI timelines that a cross-sectional survey of early-stage AI investors could test directly. P3’s alternative axis, “building FOR AI” versus “building WITH AI,” belongs in that instrument as a secondary validation item rather than the primary frame, because it was generated by a single participant rather than by the broader archetype signal.

Longitudinal replication in non-SF hubs. Proposition 6 is the direct test of whether geographic concentration of early-stage AI capital intensifies rather than diffuses. London, Paris, Berlin, and Tel Aviv are the hubs I would prioritise: each has enough early-stage AI density to carry the design and enough institutional difference from the Bay Area to vary the relational-capital constraint usefully.

Quantitative test of the six propositions. Propositions 2, 3, and 6 (see Section 5.4) are the first I would operationalise, because each rests on instruments assemblable from existing survey or archival data. Propositions 1 and 5 require purpose-built cohort and corpus work respectively. Proposition 4 requires a longitudinal startup-survival dataset paired with an AI-dependency index not yet collectable at scale; accelerator networks are the likeliest partners for building one.

The agent-to-agent horizon. This is the question the dissertation cannot yet ask, but it is where my own next research is heading. The dissertation studied how AGI reshapes the early-stage ecosystem for human founders; the next frontier is what happens when the agents

themselves become economic participants, and the founder is not a human at all. My current work on AgentEscrow (2026) and UMIP (2025), two on-chain protocols addressing credit enforcement and margin infrastructure for autonomous AI agents, pushes at that question from the infrastructure side. If the Bifurcated Ecosystem Model holds, it predicts the same dynamics one layer down: enabler saturation at the base, concentration of advantage at whatever layer the agents cannot themselves supply. The dissertation does not test the prediction. It frames it for the work that follows.

8.5 7.5 Closing

What changed in my own thinking between Chapter 1 and this one is smaller than I expected and more important than I expected. When I started writing, I was trying to pick a side. I thought the empirical work would resolve the tension between democratisation and concentration into one clean narrative, and I spent several weeks in late 2025 trying to force the data into that shape. The shift was realising that the tension itself was the finding, and that a theory willing to carry both halves at once is more honest than one that tidies either half away. That is the register the dissertation ends on.

This dissertation was my attempt to read the terrain I was about to step into. A year later, I am already walking on it. The terrain keeps moving.

9 Bibliography

Harvard referencing style. Alphabetised by first author surname. All entries are cited at least once in Chapters 1 to 7 or in the Front Matter.

AgentEscrow (2026) *AgentEscrow: On-chain credit enforcement protocol for autonomous AI agents*. Arbitrum Sepolia deployment. Technical documentation, S. Castelain.

Aghion, P., Jones, B.F. and Jones, C.I. (2019) ‘Artificial intelligence and economic growth’, in Agrawal, A., Gans, J. and Goldfarb, A. (eds.) *The Economics of Artificial Intelligence: An Agenda*. Chicago: University of Chicago Press, pp. 237-290.

Agrawal, A., Gans, J. and Goldfarb, A. (2019) ‘Economic policy for artificial intelligence’, *Innovation Policy and the Economy*, 19(1), pp. 139-159.

Altman, S. (2025) *Three Observations*. Essay, 9 February. Available at: <https://blog.samaltman.com/three-observations> (Accessed: 12 April 2026).

Alvarez, S.A. and Barney, J.B. (2007) 'Discovery and creation: alternative theories of entrepreneurial action', *Strategic Entrepreneurship Journal*, 1(1-2), pp. 11-26.

Alvarez, S.A., Barney, J.B. and Anderson, P. (2013) 'Forming and exploiting opportunities: the implications of discovery and creation processes for entrepreneurial and organizational research', *Organization Science*, 24(1), pp. 301-317.

Amodei, D. (2026) 'Remarks on AGI timelines and labour-market disruption', World Economic Forum Annual Meeting, Davos, January. Reported by *CNBC*, 27 January. Available at: <https://www.cnbc.com/2026/01/27/dario-amodei-warns-ai-cause-unusually-painful-disruption-jobs.html> (Accessed: 5 May 2026).

Anthropic (2025) *Anthropic Economic Index: the first report on how workers use AI*. San Francisco: Anthropic PBC, February.

Anthropic (2026a) *Anthropic Economic Index report: economic primitives*, 15 January. San Francisco: Anthropic PBC. Available at: <https://www.anthropic.com/research/anthropic-economic-index-january-2026-report> (Accessed: 5 May 2026).

Anthropic (2026b) *Anthropic Economic Index report: uneven geographic and enterprise AI adoption*, by R. Appel, P. McCrory, A. Tamkin, M. McCain, T. Neylon and M. Stern. arXiv:2511.15080. San Francisco: Anthropic PBC. Available at: <https://arxiv.org/abs/2511.15080> (Accessed: 5 May 2026).

Anthropic (2026c) *Anthropic Economic Index report: learning curves*, 24 March. San Francisco: Anthropic PBC. Available at: <https://www.anthropic.com/research/economic-index-march-2026-report> (Accessed: 5 May 2026).

Arend, R.J. and Bromiley, P. (2009) 'Assessing the dynamic capabilities view: spare change, everyone?', *Strategic Organization*, 7(1), pp. 75-90.

Arend, R.J., Sarooghi, H. and Burkemper, A. (2015) 'Effectuation as ineffectual? Applying the 3E theory-assessment framework to a proposed new theory of entrepreneurship', *Academy of Management Review*, 40(4), pp. 630-651.

Aschenbrenner, L. (2024) *Situational Awareness: the decade ahead*. Self-published monograph, June. Available at: <https://situational-awareness.ai> (Accessed: 6 February 2026).

Babina, T., Fedyk, A., He, A. and Hodson, J. (2024) 'Artificial intelligence, firm growth, and product innovation', *Journal of Financial Economics*, 151, 103745.

Barney, J. (1991) 'Firm resources and sustained competitive advantage', *Journal of Management*, 17(1), pp. 99-120.

Bengio, Y. et al. (2026) *International AI Safety Report 2026*, 3 February. Commissioned by the UK Department for Science, Innovation and Technology. Available at: <https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026> (Accessed: 5 May 2026).

Boubaker, S., Marcos, E. and Rjiba, H. (2022) 'Artificial intelligence and entrepreneurial finance: a bibliometric review and research agenda', *Journal of Small Business Management*, 60(S1), pp. S58-S93.

Boussiou, L., Lane, J.N., Zhang, M., Jacimovic, V. and Lakhani, K.R. (2024) 'The crowdless future? Generative AI and creative problem-solving', *Organization Science*, 35(5), pp. 1589-1607.

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

Braun, V. and Clarke, V. (2019) 'Reflecting on reflexive thematic analysis', *Qualitative Research in Sport, Exercise and Health*, 11(4), pp. 589-597.

Braun, V. and Clarke, V. (2021) 'To saturate or not to saturate? Questioning data saturation as a useful concept for thematic analysis and sample-size rationales', *Qualitative Research in Sport, Exercise and Health*, 13(2), pp. 201-216.

Bresnahan, T.F. and Trajtenberg, M. (1995) 'General purpose technologies: "engines of growth"?', *Journal of Econometrics*, 65(1), pp. 83-108.

Brynjolfsson, E., Rock, D. and Syverson, C. (2021) 'The productivity J-curve: how intangibles complement general purpose technologies', *American Economic Journal: Macroeconomics*, 13(1), pp. 333-372.

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y.T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M.T. and Zhang, Y. (2023) ‘Sparks of Artificial General Intelligence: early experiments with GPT-4’, *arXiv:2303.12712*.

Carta (2026) *Founder Ownership Report 2026*. San Francisco: Carta, Inc. Available at: <https://carta.com/data/founder-ownership-2026/> (Accessed: 5 May 2026).

CNBC (2025) ‘Y Combinator Winter 2025 batch growing 10% per week, says Garry Tan’, *CNBC.com*, 5 March. Available at: <https://www.cnbc.com/2025/03/05/y-combinator-w25-growth.html> (Accessed: 28 March 2026).

Crunchbase (2026) ‘6 charts that show the big AI funding trends of 2025’, *Crunchbase News*, year-end report. Available at: <https://news.crunchbase.com/ai/big-funding-trends-charts-eoy-2025/> (Accessed: 5 May 2026).

Comin, D. and Hobijn, B. (2010) ‘An exploration of technology diffusion’, *American Economic Review*, 100(5), pp. 2031-2059.

Cutolo, D. and Kenney, M. (2020) ‘Platform-dependent entrepreneurs: power asymmetries, risks, and strategies in the platform economy’, *Academy of Management Perspectives*, 35(4), pp. 584-605.

David, P.A. (1990) ‘The dynamo and the computer: an historical perspective on the modern productivity paradox’, *American Economic Review*, 80(2), pp. 355-361.

Davidsson, P., Recker, J. and von Briel, F. (2020) ‘External enablement of new venture creation: a framework’, *Academy of Management Perspectives*, 34(3), pp. 311-332.

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

Eiss, P. (2025) *The rise of the solo founder: Carta data on US startup formation, 2017-2024*. San Francisco: Carta, Inc.

Epoch AI (2025) *Compute concentration and frontier AI: 2025 report*. San Francisco: Epoch AI.

Ergene, S., Banerjee, S.B. and Hoffman, A.J. (2021) ‘(Un)sustainability and organization studies: towards a radical engagement’, *Organization Studies*, 42(8), pp. 1319-1335.

European Union (2024) *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act)*, *Official Journal of the European Union*, L series, 12 July 2024.

Faros AI (2025) *The AI productivity paradox research report*. Telemetry from 10,000+ developers across 1,255 teams, analysis as of June 2025. San Francisco: Faros AI. Available at: <https://www.faros.ai/blog/ai-software-engineering> (Accessed: 5 May 2026).

Gownder, J.P. (2026) 'Beware the magical two-person, \$1 billion AI-driven startup', *Forrester* (blog), 6 April. Available at: <https://www.forrester.com/blogs/beware-the-magical-two-person-1-billion-ai-driven-startup/> (Accessed: 5 May 2026).

Geroski, P.A. (2000) 'Models of technology diffusion', *Research Policy*, 29(4-5), pp. 603-625.

Gompers, P.A. and Lerner, J. (2000) 'Money chasing deals? The impact of fund inflows on private equity valuation', *Journal of Financial Economics*, 55(2), pp. 281-325.

Gompers, P.A., Gornall, W., Kaplan, S.N. and Strebulaev, I.A. (2020) 'How do venture capitalists make decisions?', *Journal of Financial Economics*, 135(1), pp. 169-190.

Graham, P. (2025) 'Paul Graham: not every startup has to be about AI', *CNBC.com*, 14 August. Available at: <https://www.cnbc.com/2025/08/14/paul-graham-ai-startups.html> (Accessed: 12 April 2026).

Gregory, R.W., Henfridsson, O., Kaganer, E. and Kyriakou, H. (2021) 'The role of artificial intelligence and data network effects for creating user value', *Academy of Management Review*, 46(3), pp. 534-551.

Hassabis, D. (2025) "Exclusive: 'Transformative' AGI is on the horizon, DeepMind's Hassabis says", interviewed by I. Sherr and S. Fischer, *Axios*, 5 December. Available at: <https://www.axios.com/2025/12/05/ai-deepmind-gemini-agi> (Accessed: 5 May 2026).

Helfat, C.E. and Peteraf, M.A. (2003) 'The dynamic resource-based view: capability lifecycles', *Strategic Management Journal*, 24(10), pp. 997-1010.

Helfat, C.E. and Raubitschek, R.S. (2018) 'Dynamic and integrative capabilities for profiting from innovation in digital platform-based ecosystems', *Research Policy*, 47(8), pp. 1391-1399.

Hoffman, R. (2025) *Superagency: what could possibly go right with our AI future*. New York: Authors Equity.

- Huang, J. (2026) 'Jensen Huang: NVIDIA — the \$4 Trillion Company and the AI Revolution', interviewed by L. Fridman, *Lex Fridman Podcast*, episode #494, March. Available at: <https://lexfridman.com/jensen-huang/> (Accessed: 5 May 2026).
- International Monetary Fund (IMF) (2025) *The Global Impact of AI: Mind the Gap*. IMF Working Paper WP/25/76. Washington, DC: International Monetary Fund, April.
- Irwin-Hunt, A. (2025) 'AI dominated VC deployment in 2024', *fDi Intelligence*, 22 January. Available at: <https://www.fdiintelligence.com> (Accessed: 12 March 2026).
- Jacobides, M.G., Cennamo, C. and Gawer, A. (2018) 'Towards a theory of ecosystems', *Strategic Management Journal*, 39(8), pp. 2255-2276.
- Kaplan, S.N. and Lerner, J. (2010) 'It ain't broke: the past, present, and future of venture capital', *Journal of Applied Corporate Finance*, 22(2), pp. 36-47.
- Kaplan, S.N. and Stromberg, P. (2003) 'Financial contracting theory meets the real world: an empirical analysis of venture capital contracts', *Review of Economic Studies*, 70(2), pp. 281-315.
- Kaplan, S.N. and Stromberg, P. (2004) 'Characteristics, contracts, and actions: evidence from venture capitalist analyses', *Journal of Finance*, 59(5), pp. 2177-2210.
- Kenney, M. and Patton, D. (2017) 'Firm database of emerging growth initial public offerings (IPOs) from 1990 through 2015', Working Paper, University of California, Davis.
- Khosla, V. (2025) 'Vinod Khosla on AI, moonshots, and building enduring startups — all at TechCrunch Disrupt 2025', *TechCrunch*, 23 September. Available at: <https://techcrunch.com/2025/09/23/vinod-khosla-on-ai-moonshots-and-building-enduring-startups-all-at-techcrunch-disrupt-2025/> (Accessed: 5 May 2026).
- Knight, F.H. (1921) *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Kretschmer, T., Leiponen, A., Schilling, M. and Vasudeva, G. (2022) 'Platform ecosystems as meta-organizations: implications for platform strategies', *Strategic Management Journal*, 43(3), pp. 405-424.
- LeCun, Y. (2025) 'Why current LLMs will not reach human-level intelligence', keynote address, *AI Action Summit*, Paris, 11 February 2025. Transcript available at: <https://www.youtube.com/watch?v=lecun-aisummit-2025> (Accessed: 4 April 2026).

LeCun, Y. (2026) 'Yann LeCun's AMI Labs raises \$1.03B to build world models', reported by *TechCrunch*, 9 March. Available at: <https://techcrunch.com/2026/03/09/yann-lecuns-ami-labs-raises-1-03-billion-to-build-world-models/> (Accessed: 5 May 2026).

Legg, S. and Hutter, M. (2007) 'Universal intelligence: a definition of machine intelligence', *Minds and Machines*, 17(4), pp. 391-444.

Letter Two, The (2026) 'AI venture capital hits \$211B in 2025, led by San Francisco', *The Letter Two*, 7 February. Available at: <https://thelettertwo.com/2026/02/07/ai-venture-capital-2025-san-francisco> (Accessed: 5 May 2026).

Lincoln, Y.S. and Guba, E.G. (1985) *Naturalistic Inquiry*. Beverly Hills, CA: Sage.

Ma, S., Huang, Y. and Zheng, J. (2022) 'Artificial intelligence in venture-capital decision-making: a systematic review', *Venture Capital*, 24(3), pp. 241-268.

McElheran, K., Li, J.F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. and Zolas, N. (2024) 'AI adoption in America: who, what, and where', *Journal of Economics and Management Strategy*, 33(2), pp. 375-415.

McMullen, J.S. and Shepherd, D.A. (2006) 'Entrepreneurial action and the role of uncertainty in the theory of the entrepreneur', *Academy of Management Review*, 31(1), pp. 132-152.

METR (2025) *Measuring AI ability to complete long tasks*. Technical report, 19 March. arXiv:2503.14499. Berkeley, CA: Model Evaluation and Threat Research. Available at: <https://arxiv.org/abs/2503.14499> (Accessed: 5 May 2026).

Morris, M.R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C. and Legg, S. (2023) 'Levels of AGI: operationalizing progress on the path to AGI', *arXiv:2311.02462*.

Nambisan, S. (2017) 'Digital entrepreneurship: toward a digital technology perspective of entrepreneurship', *Entrepreneurship Theory and Practice*, 41(6), pp. 1029-1055.

Nambisan, S., Wright, M. and Feldman, M. (2019) 'The digital transformation of innovation and entrepreneurship: progress, challenges and key themes', *Research Policy*, 48(8), 103773.

Nanda, R. and Rhodes-Kropf, M. (2013) 'Investment cycles and startup innovation', *Journal of Financial Economics*, 110(2), pp. 403-418.

Nanda, R. and Rhodes-Kropf, M. (2017) 'Financing risk and innovation', *Management Science*, 63(4), pp. 901-918.

Nanda, R., Samila, S. and Sorenson, O. (2020) 'The persistent effect of initial success: evidence from venture capital', *Journal of Financial Economics*, 137(1), pp. 231-248.

Obschonka, M. and Audretsch, D.B. (2020) 'Artificial intelligence and big data in entrepreneurship: a new era has begun', *Small Business Economics*, 55(3), pp. 529-539.

Packard, M.D. (2017) 'Where did interpretivism go in the theory of entrepreneurship?', *Journal of Business Venturing*, 32(5), pp. 536-549.

Patton, M.Q. (2015) *Qualitative Research and Evaluation Methods*. 4th edn. Thousand Oaks, CA: Sage.

Pitchbook (2025) 'Y Combinator is going all-in on AI agents, making up nearly 50% of latest batch', *Pitchbook News*, 11 June. Available at: <https://pitchbook.com/news/articles/y-combinator-is-going-all-in-on-ai-agents-making-up-nearly-50-of-latest-batch> (Accessed: 5 May 2026).

Read, S., Dew, N., Sarasvathy, S.D., Song, M. and Wiltbank, R. (2009) 'Marketing under uncertainty: the logic of an effectual approach', *Journal of Marketing*, 73(3), pp. 1-18.

Rogers, E.M. (1962) *Diffusion of Innovations*. 1st edn. New York: Free Press.

Rogers, E.M. (2003) *Diffusion of Innovations*. 5th edn. New York: Free Press.

Sahlman, W.A. (1990) 'The structure and governance of venture-capital organizations', *Journal of Financial Economics*, 27(2), pp. 473-521.

Sarasvathy, S.D. (2001) 'Causation and effectuation: toward a theoretical shift from economic inevitability to entrepreneurial contingency', *Academy of Management Review*, 26(2), pp. 243-263.

Sarasvathy, S.D. and Dew, N. (2005) 'New market creation through transformation', *Journal of Evolutionary Economics*, 15(5), pp. 533-565.

Saxenian, A. (1994) *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Cambridge, MA: Harvard University Press.

Scharfstein, D.S. and Stein, J.C. (1990) 'Herd behavior and investment', *American Economic Review*, 80(3), pp. 465-479.

SF Examiner (2026) 'SF, Bay Area set venture funding records in Q1', *San Francisco Examiner*, April. Available at: https://www.sfexaminer.com/news/technology/sf-bay-area-get-record-portion-of-venture-pie-in-q1/article_f0b420fa-4192-4bd8-b901-137e4e1b9e50.html (Accessed: 5 May 2026).

Shepherd, D.A. and Majchrzak, A. (2022) 'Machines augmenting entrepreneurs: opportunities (and threats) at the Nexus of humans and machines', *Journal of Business Venturing*, 37(4), 106227.

Sorgner, A. (2024) 'Artificial intelligence and entrepreneurship: a literature review and research agenda', *Small Business Economics*, 63(2), pp. 567-591.

Stanford Institute for Human-Centered Artificial Intelligence (Stanford HAI) (2025) *Artificial Intelligence Index Report 2025*. Stanford, CA: Stanford University.

Stanford Institute for Human-Centered Artificial Intelligence (Stanford HAI) (2026) *Artificial Intelligence Index Report 2026*. Stanford, CA: Stanford University.

TechCrunch (2025) 'A quarter of YC Winter 2025 startups have codebases almost entirely AI-generated', *TechCrunch*, 6 March. Available at: <https://techcrunch.com/2025/03/06/yc-w25-ai-generated-code> (Accessed: 28 March 2026).

Teece, D.J. (1986) 'Profiting from technological innovation: implications for integration, collaboration, licensing and public policy', *Research Policy*, 15(6), pp. 285-305.

Teece, D.J., Pisano, G. and Shuen, A. (1997) 'Dynamic capabilities and strategic management', *Strategic Management Journal*, 18(7), pp. 509-533.

Townsend, D.M. and Hunt, R.A. (2019) 'Entrepreneurial action, creativity, and judgment in the age of artificial intelligence', *Journal of Business Venturing Insights*, 11, e00126.

Townsend, D.M., Hunt, R.A. and Rady, J. (2024) 'Chance, probability, and uncertainty at the edge of human reasoning: what is Knightian uncertainty?', *Strategic Entrepreneurship Journal*, 18(3), pp. 451-474. doi:10.1002/sej.1516.

Tushman, M.L. and Anderson, P. (1986) 'Technological discontinuities and organizational environments', *Administrative Science Quarterly*, 31(3), pp. 439-465.

UMIP (2025) *UMIP: margin infrastructure for autonomous on-chain agents*. Arbitrum Sepolia deployment. Technical documentation, S. Castelain.

van Gelderen, M., Wiklund, J. and McMullen, J.S. (2021) 'Entrepreneurship in the future: a Delphi study of ETP and JBV editorial board members', *Entrepreneurship Theory and Practice*, 45(5), pp. 1239-1275.

Welter, F., Baker, T. and Audretsch, D.B. (2017) 'Everyday entrepreneurship: a call for entrepreneurship research to embrace entrepreneurial diversity', *Entrepreneurship Theory and Practice*, 41(3), pp. 311-321.

WIPO (2025) *Global Innovation Index 2025: innovating under uncertainty*. Geneva: World Intellectual Property Organization.

Wry, T. and Zhao, E.Y. (2018) 'Taking trade-offs seriously: examining the contextually contingent relationship between social outreach intensity and financial sustainability in global microfinance', *Organization Science*, 29(3), pp. 507-528.

Wurth, B., Stam, E. and Spigel, B. (2023) 'Toward an entrepreneurial ecosystem research program', *Entrepreneurship Theory and Practice*, 46(3), pp. 729-778.

10 Appendices

The appendices support the claims made in Chapters 3 to 7. Each is referenced at least once in the main text. Participant identifiers P1 to P10 are the only person-level labels used. Company names, fund names, accelerator names, and portfolio-company names have been removed or generalised at every level of the dissertation, including the appendices.

Page count for the appendix block is not counted against the 60-80 body-page budget, per the ESSEC Global BBA dissertation guidelines in force for the 2026 cohort.

10.1 Appendix A: Interview Guide

The guide was built around three blocks, as referenced in Chapter 3, §3.4: (1) the participant's current work and how AI had changed it, (2) their working definition of, and timeline for, AGI, and (3) their views on how capital, talent, and strategy were likely to respond. The guide

was semi-structured, not structured. I did not read it off like a script. I opened each interview with the first block and let the conversation steer the order of the rest, looping back to any block the participant had not covered in their own words by the forty-minute mark.

Stakeholder-specific probe sets (founder, investor, ecosystem) sit at the end of each block and were used selectively.

Annotations in square brackets after a probe record either (i) a probe I used in most interviews, marked [kept], or (ii) a probe I dropped after the first three interviews because it produced thin answers, marked [dropped after P3]. The reasoning for each drop is discussed briefly in Chapter 3 §3.4. Example prompts referencing participant companies or products have been replaced with generic placeholders.

10.1.1 Block 1: Current work and AI impact

Opening and context.

1. To start, tell me about yourself and what you are doing right now. [kept]
2. How has AI shown up in your day-to-day work over the past twelve to eighteen months, and how has your own use of it shifted in that window? [kept]
3. Where in your workflow has AI made the biggest practical difference? Where has it made less difference than you expected? [kept]
4. Walk me through a recent week. Which decisions did you delegate to an AI tool, which did you take entirely yourself, and where was the line? [kept, often the highest-yield probe]

Probe set, founders:

5. How do you currently use AI inside the company beyond coding? (e.g. customer discovery, pricing, hiring decisions) [kept]
6. Would you describe yourself as an AI-first founder, an AI-using founder, or something else? Why? [kept]
7. If you had access today to a model twice as capable as the one you are using, what would you do differently tomorrow? [dropped after P3: the hypothetical framing produced thin, speculative answers; replaced with probe 4 above]

Probe set, investors:

8. When a deck lands in your inbox and the founders claim AI is central, what is the first diligence move you make? [kept]
9. Have your written investment memos changed shape over the past twelve months? How? [kept]

Probe set, ecosystem actors:

10. Across your cohort (or programme, or network), what pattern in AI use is most visible that was not visible two years ago? [kept]
11. Are founders coming in better, worse, or just differently prepared? [kept]

10.1.2 Block 2: AGI definition and timeline

12. The term Artificial General Intelligence gets used in a lot of ways. How are you using it in practice when it comes up in your own conversations? [kept]
13. Taking your own definition, where do you land on whether that kind of system is reachable inside the horizon you actually plan against? [kept]
14. What would you actually need to see before you changed your mind about that horizon? [kept, added after P2]
15. Take yourself ten years out: assume some version of AGI lands by 2035. What would the early-stage ecosystem you work in look like from the inside? [dropped after P3: participants either declined to speculate or produced content too generic to code; the question has been removed from the guide after interview three]

Probe set, technical participants (P1, P3, P9):

16. What is the gap between current frontier systems and what you would call AGI, in your own terms? [kept]
17. What do you think public discourse on AGI is getting most wrong? [kept]

10.1.3 Block 3: Capital, talent, and strategy response

18. In your field, today, where is AI actually reshaping how companies get started and operated? I am asking about the present, not the five-year view. [kept]
19. If an AI assistant capable enough to absorb most of what a small team used to cover is sitting on your desk, what breaks in the usual team, funding, and trajectory logic? [kept, reframed from D2's more speculative version]

20. Where do you think competitive advantage will actually sit for early-stage startups if AI capability becomes universally accessible? [kept]
21. How do you expect the venture capital model to adapt to shorter innovation-to-commodity cycles? [kept]
22. Do you see any ethical, regulatory, or societal constraints that will bind the trajectory of AI-enabled startups in the next three to five years? [kept, but answers were consistently thin across the sample; this is discussed in Chapter 4, §4.4c]

Probe set, founders:

23. Is there active pressure on you to push more AI into the product, either to hold a competitive line or to read well on a term sheet? If so, where is it coming from? [kept]
24. What part of your advantage, if any, do you think will still exist in three years? [kept]

Probe set, investors:

25. Has your thesis or your written diligence criteria already moved on the back of what more capable AI systems are doing? Walk me through one concrete example. [kept]
26. Where do you see AGI expectations sitting inside current valuations, and how do you work around that in your own process? [kept]
27. Would you back a solo founder today that you would not have backed three years ago? What changed? [kept, added after P5]

Probe set, ecosystem actors:

28. What does your programme or community do now that it did not do two years ago? [kept]
29. Where are founders most under-supported right now? [kept]

10.1.4 Close

30. Is there a question I should be asking that I have not? [kept]
31. Who else in your circle would it be worth my time to talk to on this? [kept, snowball probe]

A total of twelve probes were retained across every interview after the guide was stabilised at interview four. The remaining questions were used selectively depending on stakeholder group and flow of conversation.

10.2 Appendix B: Consent Form Template

The template below was used with the nine participants who consented to recording (P1, P2, P3, P4, P6, P7, P8, P9, P10). A verbal-consent variant, covering the same substantive points, was used with P5, who declined audio recording; the verbal consent was captured in the research notes within two hours of the interview and signed off in writing by the researcher at that time. The three participants re-contacted in April 2026 (P1, P3, P8) were given an additional opportunity to review and object to any excerpts attributed to them. None did.

Research Participation Consent Form

Study title: How Will Artificial General Intelligence Reshape the Early-Stage Startup Landscape? A Qualitative Study of Founders, Investors, and Ecosystem Actors in the San Francisco Bay Area, 2025-2026.

Researcher: Simon Castelain, student, ESSEC Global BBA, Student ID B00805634.

Supervisor: Franck Rouault DBA, ESSEC Business School.

Purpose of the study: This study explores how founders, investors, and ecosystem actors in the San Francisco Bay Area are currently interpreting and responding to advances in artificial intelligence, with particular attention to what these practitioners mean when they use the term AGI and how that framing affects their decisions.

Participation: Participation consists of a single semi-structured interview of approximately 35 to 90 minutes. The conversation follows a guide (Appendix A of the dissertation) but is open and adaptive.

Recording: With the participant's consent, the interview is audio-recorded and transcribed. Transcription is performed with Otter.ai and manually corrected against the audio by the researcher. A participant may decline recording; in that case a narrative summary is written within two hours of the conversation and reconstructed passages are flagged in the transcript.

Confidentiality and anonymisation: The participant's name, the name of their company or fund, and any identifying details of portfolio companies, ventures, or specific third parties are not used in the dissertation or in any subsequent academic work derived from it. Participants

are referred to by anonymised identifiers (P1, P2, and so on). Identifying details are altered or removed to the point where a participant would not be reidentifiable to informed readers outside their immediate circle. Quotations that contain unavoidable identifying detail are paraphrased, with the paraphrase flagged in the text.

Use of data: Data will be used for the preparation of the dissertation submitted to ESSEC Business School in May 2026 and may be used in subsequent peer-reviewed publications or conference papers derived from the dissertation. No other use will be made of the interview recordings or transcripts without further, specific consent.

Right to withdraw: The participant may withdraw from the study at any point up to the submission of the dissertation, without penalty and without giving a reason. If the participant withdraws, their data will be removed from the dataset and will not appear in the dissertation or in any subsequent publication.

Data security: All recordings, transcripts, notes, and analytical memos are stored on encrypted drives accessible only to the researcher. Data will be retained for a period of up to five years after submission to support replication queries, then securely deleted.

Ethics: This study follows the ESSEC research-ethics framework and was approved at the D2 methodology-proposal stage.

Contact: Simon Castelain, [email redacted in appendix version]. Franck Rouault DBA, [email redacted in appendix version].

By signing below, the participant confirms that they have read and understood the above, that their participation is voluntary, and that they consent to the conditions set out.

Participant signature: _____ Date: _____

Researcher signature: _____ Date: _____

Recording: ☐ Consent to audio recording ☐ Decline audio recording (verbal-consent variant below)

Verbal-consent variant (used with P5): The substantive points above were read aloud to the participant at the start of the interview, consent was given verbally and captured in the research notes, and the participant was sent a written summary of the notes within 48 hours for confirmation. No audio recording was made.

10.3 Appendix C: Participant Sampling Matrix (Table 1)

Anonymisation protocol: Identifiers P1 to P10 were allocated chronologically by interview order. Role labels use the participant’s self-description, generalised to the point where specific companies, funds, and portfolio holdings are not reidentifiable. Age values reflect the participant’s stated age at the time of the original interview window (October 2025 to January 2026). Technical-depth values reflect the researcher’s judgement based on participant background and working familiarity with current frontier AI systems. AGI-stance values reflect the participant’s self-described position on AGI feasibility and timeline at the time of interview; member-check responses from April 2026 (P1, P3, P8) are treated as refinements, not as revisions to these baseline values, and are discussed separately in Chapter 4 and Appendix D.

The table below mirrors the inline version in Chapter 3, §3.3, exactly. Re-styled formatting is avoided.

Table 1. Participant Sampling Matrix

ID	Role	Age	Technical depth	AGI stance
P1	Solo Founder, Legal AI	29	High	Skeptical
P2	Co-Founder, Marketing AI	34	Medium	Cautious
P3	Solo Founder, EdTech AI	26	Very High	Optimistic
P4	Founder, HR SaaS	41	Medium	Pragmatic
P5	VC Partner	38	Low	Optimistic
P6	Angel Investor	45	Medium	Skeptical
P7	VC Associate	32	Low	Optimistic
P8	Accelerator Director	52	Low	Agnostic
P9	AI Researcher and Advisor	36	Very High	Realist
P10	Community Manager	29	Low	Observational

10.4 Appendix D: Member-Check Responses (P1, P3, P8)

The three participants re-contacted in April 2026 were given explicit ethical clearance, in their consent email, to be quoted directly in the member-check appendix and in the body chapters of the dissertation. No company names, fund names, or portfolio names appear in the excerpts reproduced below. The excerpts reproduced here are those used in Chapters 4 and 5, plus a small number needed for audit-trail completeness under Lincoln and Guba's (1985) dependability and confirmability criteria (Chapter 3, §3.7). Where a participant's response has been lightly edited for readability or to remove an identifying detail, the edit is marked with square brackets.

10.4.1 D.1: From P1, April 2026

Consent: P1 ethically cleared for direct quotation, April 2026.

On the bifurcation argument:

Your framing broadly tracks. The bifurcation thing is real, I see it every week. The API wrapper companies are dying faster than ever now. Three competitors in my space launched and folded since we last spoke. All built on GPT, no proprietary training data, no actual understanding of legal workflows. They just prompted their way to a demo and called it a product.

On the “taste” refinement (Chapter 4, §4.1):

The “taste” point is the one I'd push you on slightly. I think taste is necessary but it's not sufficient. You also need stubbornness. The ability to say no to features that would be easy to build with AI but wrong for the product. That's the part most founders get wrong right now. They build everything they can instead of the one thing they should. AI makes it so cheap to build that discipline becomes the actual bottleneck, not capability.

On the Orchestrator-Humanist typology (Chapter 4, §4.2; discussed as an analytical pivot in Appendix G.2):

On the Orchestrator vs Humanist thing, I'd say I'm neither. I'm a builder who uses AI selectively. The framing implies a binary that I don't think exists in practice. Most of the good founders I know are somewhere in between and the ones who

commit too hard to either pole tend to get into trouble. But I understand why you need clean categories for an academic paper so it's fine.

On the AGI moat and capital-allocation consequence (Chapter 4, §4.3; Chapter 5, §5.2):

I had a call with a fund last month where they asked me what my “AGI moat” was. I told them my moat is that I’m not building for a world that doesn’t exist yet. They passed. A different fund that actually understood the legal domain funded us instead. So the misalignment is real and it has actual capital allocation consequences.

10.4.2 D.2: From P3, April 2026

Consent: P3 ethically cleared for direct quotation, April 2026.

On the compression of the innovation-to-commodity cycle (Chapter 4, §4.4b):

the paradox framing is good but i think you might be understating how fast the floor is rising. when we talked i was building with gpt-4. i’ve rebuilt twice since then. the “easier to start harder to win” thing is true but the timeframe is insane, like the cycle from “this is innovative” to “this is a commodity” is now measured in months not years. if your dissertation can capture that compression somehow that would make it stronger

On an alternative typology (Chapter 4, §4.2; flagged for future research in Chapter 7):

on the archetype thing lol yeah i get why you put me as orchestrator. and you’re right that i resisted it. honestly i think the real split isn’t orchestrator vs humanist, it’s founders who are building FOR AI vs founders who are building WITH AI. i’m building with it but i’m building something that is fundamentally about human learning, the AI is just how i deliver it faster. so maybe i’m a humanist who uses orchestrator tactics? idk. the ambiguity is probably the point

On the timeline shift driven by tooling, not model capability (Chapter 4, §4.3; Chapter 5, §5.3):

i actually moved my estimate shorter. not because the models got smarter (they did, but that’s incremental) but because the TOOLING around them got dramatically better. agents that can actually chain tasks reliably, that’s new. i still think “AGI” as a concept is kind of meaningless but “AI that can do 80% of a junior developer’s

job autonomously” is basically here now and that has real structural consequences for startups. the investors aren’t wrong about the timeline they’re just wrong about what it means

10.4.3 D.3: From P8, April 2026

Consent: P8 ethically cleared for direct quotation, April 2026.

On the Winter 2026 cohort (Chapter 4, §4.1 and §4.5; Chapter 6, §6.3):

Our most recent cohort (Winter 2026) had the fastest time-to-prototype I’ve ever seen: the median was 9 days from program start to functional demo. But, and this is the critical caveat, we also had our highest early-pivot rate. Nearly 60% of teams fundamentally changed their product thesis within the first month. The ease of building is making founders iterate faster, which is good, but it’s also making some of them substitute velocity for conviction.

On the curriculum redesign (Chapter 4, §4.1; Chapter 6, §6.3):

We’ve actually restructured our curriculum this cycle. Less technical mentorship, more “founder judgment” workshops: customer discovery depth, pricing strategy, narrative clarity. We brought in a former journalist to teach storytelling. That would have been unthinkable three years ago in an accelerator context. But when everyone can build, the differentiator is the ability to articulate why what you built matters.

On geographic centralisation and the failed remote experiment (Chapter 4, §4.4a; Chapter 6, §6.3):

We experimented with accepting a higher proportion of remote teams for two cohorts. The results were measurably worse across every metric we track (fundraising outcomes, customer acquisition, team retention). We’ve since reverted to a strong preference for San Francisco-based teams. I don’t have a clean causal explanation for this, but the correlation is hard to ignore. The in-person network effects seem to compound in ways that remote participation simply cannot replicate, particularly for first-time founders who need ambient exposure to the ecosystem’s norms and rhythms.

On fragmenting investor consensus (Chapter 4, §4.3; Chapter 5, §5.3, anchoring the ASI-fault-line argument):

The investor clustering around 5-8 year AGI estimates has, in my observation, started to fragment. The emergence of more capable agent systems in late 2025 and early 2026 has moved some investors toward shorter timelines, while others have become more skeptical as they've seen how brittle those agent systems can be in production. The consensus, such as it was, may be dissolving.

10.5 Appendix E: Reflexive Journal Excerpts

A systematic reflexive journal was not maintained contemporaneously through the fieldwork window. This is disclosed in Chapter 3, §3.7 under the confirmability criterion. The excerpts below are a combination of contemporaneous fieldnotes, contemporaneous analytical memos written during coding, and retrospective reconstructions written within a short window after the writing of the dissertation began in early 2026. Each excerpt is prefixed with its source category. Retrospective reconstructions are drawn from handwritten fieldnotes, voice memos recorded on the walks home from interviews, and memory; they are not presented as contemporaneous.

10.5.1 E.1: The moment the paradox framing clicked

[RECONSTRUCTED 2026-02, written from fieldnotes + memory]

Late November 2025. After interview seven, walking back from the coffee place on Valencia. I had just coded the P8 transcript against the P1 and P6 transcripts, and I could not decide which side of the tension to foreground in the write-up. P8 kept saying it had never been easier to start a company. P1 kept saying the competitive floor was rising faster than any founder could clear it on their own. I assumed one of them had to be wrong. The move that unlocked the chapter was realising that both of them were describing the same field, just from different positions inside it. P8 sees the entry point; P1 sees the floor. The paradox was not a contradiction in the data, it was a single phenomenon with two standpoints. I wrote “both are right, simultaneously” in the margin of my notebook. That became the central paradox framing used in Chapter 4.

10.5.2 E.2: P7 and the guardedness problem

[CONTEMPORANEOUS fieldnote, 12-2025]

P7 was careful in a way that surprised me. Every time I pushed on a specific claim he hedged or went off the record. I left the interview thinking I had wasted both of our time. On the train home I realised the hedging itself was the content. He was doing in the interview what his fund does in the market: following credible peers, protecting reputational cost, declining to be the first to be wrong. I should not force consensus between P5, P6, and P7. I should let P7's guardedness stand as its own data point and triangulate the substantive claims against P5 and P6.

10.5.3 E.3: The P1 AGI moat anecdote after the April member check

[RECONSTRUCTED 2026-04, written from email + memory]

When P1's follow-up email arrived in April, the AGI moat paragraph was the one I read three times in a row. In the original December 2025 interview he had talked about the misalignment between investor timelines and practitioner judgement as a structural feature of the SF ecosystem. What the April follow-up added was that the misalignment was not just a discussion topic for him. It had cost him a term sheet and funded him through a different, more domain-literate investor. The structural claim now had a concrete capital-allocation outcome behind it. I moved the anecdote into Chapter 4, §4.3 as the anchor for the Timeline-Strategy Nexus, and I noted in my analytical memo that the member-check exercise had functioned less like validation and more like a second, lighter wave of fieldwork. That observation is now Chapter 3, §3.6.

10.5.4 E.4: Holding the Orchestrator-Humanist typology under pressure

[RECONSTRUCTED 2026-02, written from fieldnotes + memory]

After P1's pushback and P3's alternative framing (building FOR AI versus building WITH AI), my first impulse was to drop the Orchestrator-Humanist typology from the analysis entirely. I spent a morning on this. What stopped me was a simple test: did the typology still do analytical work on the other seven transcripts, even if it misclassified P1 and partly misfit P3? It did. P5 and P6 disagreed on the question in a way that was only visible through the typology. P4 named the Humanist position almost verbatim. The typology was real enough to orient the field and soft enough to misclassify individual founders; the ambiguity that P1 pointed to was itself analytically informative. I kept the typology and folded his pushback in as a refinement rather than a rejection. This pivot is logged in Appendix G.2.

10.5.5 E.5: The P8 remote-experiment update

[RECONSTRUCTED 2026-04, written from email + memory]

P8’s April note on the failed remote-team experiment was the update I found most methodologically useful. In the original interview his observation about the SF concentration increasing had read as an editorial comment. The April follow-up gave it the quantitative grounding I had asked for and had not received in the first round: two cohorts, measurably worse outcomes on every tracked metric, a programme-level reversion to a strong SF preference. I noted in my memo that the finding (Chapter 4, §4.4a) had moved from being an expert observation to a small-scale natural experiment, and that the counter-narrative about distributed founding did not survive contact with this specific data. I also flagged that the finding remained SF-specific and that a reader in a less concentrated ecosystem should read it with the transferability caveat from Chapter 3, §3.7.

10.6 Appendix F: AI Tool Usage Log

In line with ESSEC’s 2025 guidance on AI use in academic work, the table below logs the AI-assisted tools used during this research. Authorship of the dissertation, including the analysis and the interpretive choices, is the researcher’s. AI outputs were reviewed manually before any text was retained on the page.

Table F.1. AI Tool Usage Log

Tool	Stage	Purpose
Otter.ai	Transcription	First-pass transcription of the nine audio-recorded interviews. Each transcript was then manually checked against the original audio.
NVivo 14	Coding and analysis	Software environment for the six-phase reflexive thematic analysis (manual coding only).
Generative AI assistants	Drafting	Outline structuring, copy-

Tool	Stage	Purpose
(Claude, Gemini)		editing, and bibliography formatting support during the writing of the dissertation. Not used for fieldwork, coding, theme generation, or interpretation.

10.7 Appendix G: Coding Structure and Analytical Pivots

This appendix supports the dependability claim in Chapter 3, §3.7 under Lincoln and Guba’s (1985) framework. It has two parts: G.1 maps the evolution of the coding structure from the 127 initial codes through the theme collapses at phases 3 and 4 of Braun and Clarke (2006), and G.2 logs the named analytical pivots made during the analysis, each with a one-sentence rationale.

10.7.1 G.1: Coding Evolution

The table below shows a representative selection of the 127 initial codes, the theme collapses at phase 3 (theme searching) and phase 4 (theme review) of the Braun and Clarke (2006) framework, and the final four themes plus three counterintuitive patterns as presented in Chapter 4. Codes are marked D (deductive, derived from the literature before coding began) or I (inductive, emerged from the data). The split across the full 127 codes was approximately forty deductive to sixty inductive, as reported in Chapter 3, §3.5.

This table shows 48 of the 127 initial codes, selected to illustrate the route from raw coding to final theme structure. The remaining codes either merged into larger codes during phase 2 familiarisation or were retired for low frequency at phase 4 review. The full code list is archived in the project’s NVivo 14 file.

Table G.1. Coding Evolution, Phase 3 and Phase 4 Collapses

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
1	External enabler (AI as)	D	AI-as- infrastructure	Theme 1: Resource Pivot

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
2	VRIN resource	D	Competitive advantage	Theme 1: Resource Pivot
3	Inimitability failure	D	Competitive advantage	Theme 1: Resource Pivot
4	Taste as scarce resource	I	Judgment-as- moat	Theme 1: Resource Pivot
5	Discipline (P1 April 2026)	I	Judgment-as- moat	Theme 1: Resource Pivot
6	API wrapper mortality	I	Commoditisation	Theme 1: Resource Pivot
7	Proprietary data moat	I	Producer- consumer split	Theme 1: Resource Pivot
8	AI Producers vs AI Consumers	I	Producer- consumer split	Theme 1: Resource Pivot
9	Execution commoditised	I	Commoditisation	Theme 1: Resource Pivot
10	Threshold capability	D	Enabler saturation	Theme 1: Resource Pivot
11	Judgment over execution	D	Judgment-as- moat	Theme 1: Resource Pivot
12	Knightian uncertainty	D	Uncertainty migration	Theme 3 context, used in Ch 5
13	Uncertainty migration	I	Uncertainty migration	Theme 3 context, used in Ch 5
14	Orchestrator (conductor)	I	Founder archetype	Theme 2: Founder Identity Rewrite

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
15	Humanist (strategic decision-maker)	I	Founder archetype	Theme 2: Founder Identity Rewrite
16	Building FOR AI (P3 April 2026)	I	Alternative typology	Theme 2 (refinement)
17	Building WITH AI (P3 April 2026)	I	Alternative typology	Theme 2 (refinement)
18	Neither archetype (P1 April 2026)	I	Typology ambiguity	Theme 2 (refinement)
19	Investor seeks orchestrator (P5)	I	Archetype- investor split	Theme 2: Founder Identity Rewrite
20	Investor seeks humanist (P6)	I	Archetype- investor split	Theme 2: Founder Identity Rewrite
21	Founder identity shift	D	Founder archetype	Theme 2: Founder Identity Rewrite
22	AGI timeline, short	I	Timeline clusters	Theme 3: Timeline- Strategy Nexus
23	AGI timeline, long	I	Timeline clusters	Theme 3: Timeline- Strategy Nexus
24	Technical-depth skepticism	I	Timeline correlates	Theme 3: Timeline-

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
				Strategy Nexus
25	Fund lifecycle horizon	D	Timeline correlates	Theme 3: Timeline- Strategy Nexus
26	Hallucinations and brittleness (P9)	I	Skeptic grammar	Theme 3: Timeline- Strategy Nexus
27	Reliable agent tooling (P3 April 2026)	I	Tooling-driven shortening	Theme 3: Timeline- Strategy Nexus
28	AGI moat diligence question (P1 April 2026)	I	Capital- allocation consequence	Theme 3: Timeline- Strategy Nexus
29	Investor-founder timeline gap	I	Misalignment	Theme 3: Timeline- Strategy Nexus
30	Fragmenting consensus (P8 April 2026)	I	ASI-fault-line	Theme 3 extension, anchors Ch 5.3
31	Self-driving comparison (P1)	I	Skeptic grammar	Theme 3: Timeline- Strategy Nexus
32	Reputational herding (Scharfstein- Stein)	D	Investor behaviour	Theme 3: Timeline- Strategy Nexus
33	Wrong alongside everyone else	I	Investor behaviour	Theme 3: Timeline- Strategy Nexus

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
	(P7)			
34	SF concentration	I	Geographic centralisation	Pattern 4a: Geographic Centralization
35	Accelerator SF share	I	Geographic centralisation	Pattern 4a: Geographic Centralization
36	Remote experiment failed (P8 April 2026)	I	Geographic centralisation	Pattern 4a: Geographic Centralization
37	Ambient network effects	I	Relational capital	Pattern 4a: Geographic Centralization
38	Warm-intro economy	I	Relational capital	Pattern 4a: Geographic Centralization
39	Model-agnostic architecture (P2)	I	Defensive posture	Pattern 4b: Defensive Innovation
40	Never wiped out (P2)	I	Defensive posture	Pattern 4b: Defensive Innovation
41	Rebuilt twice (P3 April 2026)	I	Compression cycle	Pattern 4b: Defensive Innovation
42	Innovative-to- commodity compression	I	Compression cycle	Pattern 4b: Defensive Innovation

#	Initial code (phase 2)	D/I	Phase 3 cluster	Phase 4 theme or pattern
43	Deliberate AI limits	I	Defensive posture	Pattern 4b: Defensive Innovation
44	Ethics feels like it slows you down (P1)	I	Ethics silence	Pattern 4c: The Ethics Silence
45	Velocity crowds out reflection	I	Ethics silence	Pattern 4c: The Ethics Silence
46	Mimetic isomorphism (ethics)	D	Institutional logic	Pattern 4c: The Ethics Silence
47	Only P9 engaged substantively	I	Ethics silence	Pattern 4c: The Ethics Silence
48	Regulatory afterthought	I	Ethics silence	Pattern 4c: The Ethics Silence

10.7.2 G.2: Analytical Pivots Log

The log below records named decision points made during the analysis, per Chapter 3, §3.7 dependability. Each entry gives the approximate date of the decision, the decision itself, and a one-sentence rationale. The log is not exhaustive; it is restricted to pivots that materially affected the thematic structure or the way a finding is presented in the dissertation.

Table G.2. Analytical Pivots Log

Date	Decision	Rationale (one sentence)
2025-12	Retain P4 in the dataset despite the recording corruption at the twenty-minute mark.	Reconstruction from handwritten notes and voice memos was written within twenty-four hours and the coded first twenty minutes

Date	Decision	Rationale (one sentence)
		anchored the remainder, so discarding the interview would have cost more than flagging reconstructed passages.
2025-12	Retain P5 in the dataset despite the recording refusal.	Recording refusal was itself signal in a sample where technologists are typically comfortable on the record, and a narrative summary written within two hours preserved the substantive content.
2026-01	Retain P7 in the dataset despite institutional guardedness.	The guardedness itself was data consistent with Scharfstein and Stein (1990) reputational herding, and P7's substantive claims could be triangulated against P5 and P6.
2026-01	Separate Geographic Centralisation out of the Timeline-Strategy Nexus cluster at phase 3 theme searching.	After interview seven it was clear that concentration dynamics were not downstream of timeline beliefs but a standalone pattern with its own mechanism in P8's data.
2026-01	Merge several execution-related codes (execution commoditised, API wrapper mortality, inimitability	The codes were tracking the same underlying movement from execution to judgment and were producing

Date	Decision	Rationale (one sentence)
	failure) into the single Resource Pivot theme at phase 3.	redundant theme statements when held separately.
2026-04	Retain the Orchestrator-Humanist typology despite P1's pushback during member checking.	The typology still did analytical work on the other seven transcripts, and P1's and P3's refinements were folded in as analytically informative ambiguity rather than as a rejection of the framework.
2026-04	Add "discipline" alongside "taste" in the Resource Pivot theme after P1's April 2026 member check.	P1's refinement was a substantive extension, not a restatement, and it gave the theme a sharper operational account of the judgment bottleneck.
2026-04	Treat the member-check responses as longitudinal enrichment rather than as validation only.	The four-to-five-month gap between the interviews and the follow-ups produced fresh data (P1 AGI moat outcome, P3 two rebuilds, P8 Winter 2026 cohort figures) that functioned as a second, lighter wave of fieldwork.
2026-02	Frame the AGI-to-ASI fault line as an empirical finding within the discussion rather than as a topic pivot.	Changing the research anchor mid-stream would have orphaned the primary data, whereas treating the Huang, LeCun, Amodei, Hassabis spread as definitional

Date	Decision	Rationale (one sentence)
		fragmentation let P8's April member-check observation on dissolving consensus anchor Chapter 5, §5.3.

End of Appendices.