

Who Needs Wall Street? The Role of Private Capital in the Adoption of Dual-Class Share Structures

Raffaele Felicetti

SUPERVISOR

284261

CANDIDATE

Statement of Originality

I hereby declare that this thesis is my own original work completed under the guidance of my supervisor. Except where sources are properly cited, this thesis does not contain any material previously published or written by another person or organization. All individuals and groups who contributed substantially to this work have been clearly acknowledged.

Student Signature: 

Date: 05/08/2026

Statement of Authorization for Use

I fully understand the university's regulations regarding the submission, storage, and use of thesis materials. I agree to submit both printed and electronic copies of this thesis. I also authorize the university to preserve and reproduce all or part of this thesis by photocopying, digital archiving, or other appropriate means.

Student Signature: 

Supervisor Signature(s): 

Date: 05/08/2026

AI Use Statement

Artificial intelligence (AI) tools were used only in a limited, supplementary capacity during the preparation of this thesis, primarily for language refinement and organizational support. No AI tools were used to generate original arguments, conduct analysis, or produce substantive academic content.

All research design, theoretical framing, data analysis, and conclusions presented in this thesis are entirely my own work. Any AI-assisted outputs were carefully reviewed, revised, and verified for accuracy and appropriateness.

I accept full responsibility for the integrity, originality, and accuracy of this thesis.

Student Signature: 

Supervisor Signature(s): 

Date: 05/08/2026

Abstract

The empirical literature on dual-class structures in U.S. IPOs traces their adoption primarily to firm-level determinants, leaving underexplored the role of macro-financial conditions in the private capital market that may be associated with founders' bargaining power at the moment of listing and, within that environment, the distinction between potential capital (dry powder) and realized investment activity (deal volume). This study estimates, by logistic regression on 3,199 operating-company U.S. IPOs over 2000 through 2023, the association between aggregate annual late-stage venture capital deal volume and the probability of dual-class adoption, with year-clustered standard errors, 2-digit SIC industry fixed effects, and a linear time trend. The coefficient on late-stage deal volume is positive and statistically significant ($\beta = 0.0038$, $p < 0.001$ asymptotic; $p = 0.014$ under wild cluster bootstrap); when deal volume and dry powder enter jointly, deal volume retains significance at the 5% level while dry powder attenuates to marginal significance, a pattern consistent with a bargaining power channel operating through realized deal-making rather than through the latent stock of uncommitted capital. The association survives alternative measurement, functional form, profitability controls, subsample restrictions, and a wild cluster bootstrap appropriate to the 24-cluster regime; the design is associative, not causal. The thesis interprets realized late-stage deal-making as a candidate upstream mechanism whose operation is consistent with the observed pattern, situated outside the domain of the public-market instruments built to discipline a Berle-Means separation of ownership and control whose structural form the moment of listing now systematically inverts.

Keyword: dual-class IPO, late-stage deal volume, founder bargaining power, private capital market, U.S.

IPO governance

Table of Contents

Abstract.....	III
Table of Contents	IV
List of Figures and Tables	VI
1 Introduction.....	1
2 Literature Review, Theoretical Framework, and Hypothesis	4
2.1 The structural transformation of the U.S. private capital	4
2.2 The adoption of dual-class structures at the moment of the IPO	5
2.3 Founder vision and agency costs: the debate on efficiency	7
2.4 Research gap, conceptual framework and development of the hypothesis.....	9
3 Methodology	12
3.1 Research design.....	12
3.2 Unit of analysis.....	12
3.3 Sample construction and data sources.....	13
3.3.1 Population and time frame.....	13
3.3.2 Inclusion and exclusion criteria	13
3.3.3 Data sources.....	14
3.3.4 Dataset integration.....	14
3.3.5 Final sample.....	15
3.3.6 Replicability and data access	15
3.4 Variable operationalization	15
3.4.1 Dependent variable	15
3.4.2 Independent variable.....	16
3.4.3 Control variables.....	16
3.5 Model specification and standard errors	18
3.5.1 Model specification	18
3.5.2 Standard errors and clustering	19
3.6 Robustness and sensitivity analyses	19
3.7 Validity, reliability, and limitations	20

3.7.1	Internal validity.....	20
3.7.2	Construct validity.....	21
3.7.3	External validity.....	22
3.7.4	Reliability	22
3.7.5	Scope conditions and limitations	22
4	Empirical Results	24
4.1	Descriptive statistics and group comparisons	24
4.2	Main results: four specifications	25
4.3	Economic magnitude and marginal effects	26
4.4	Dry powder versus deal volume	27
4.5	Diagnostics	27
4.6	Robustness and sensitivity.....	28
5	Discussion.....	30
5.1	Interpretation of the main result	30
5.2	Deal volume and dry powder: measurement as methodological contribution	32
5.3	Robustness: the structure of the evidence	33
5.4	Normative and market implications	34
5.4.1	The inversion of the Berle-Means paradigm	34
5.4.2	The temporal asymmetry of the governance choice	36
5.4.3	The recomposition of the investable universe and the limits of index exclusion.....	37
5.4.4	Dual-class adoption at the tipping point	37
5.5	Limitations.....	38
5.6	Directions for future research.....	40
6	Conclusion	42
	References	44
	Appendix	46
	Acknowledgements.....	56

List of Figures and Tables

Figures

Figure A1: <i>The Bargaining Power Channel From Late-Stage Private Capital to Dual-Class Adoption at IPO</i>	46
---	----

Tables

Table A1: <i>Sample Attrition Process</i>	47
Table A2: <i>Variable Summary</i>	48
Table A3: <i>Descriptive Statistics (N = 3,250)</i>	49
Table A4: <i>Logistic Regression Results - Progressive Specifications (Standard Errors Clustered by Year)</i>	50
Table A5: <i>Average Marginal Effects - Model 4 (Standard Errors Clustered by Year)</i>	51
Table A6: <i>DryPowder_t1 versus LateStageVol - Full Specification (Standard Errors Clustered by Year)</i>	52
Table A7a: <i>Within-Design Robustness (Standard Errors Clustered by Year)</i>	53
Table A7b: <i>Structural Robustness - Alternative Time Trends (Standard Errors Clustered by Year)</i>	53
Table A8: <i>Inferential Diagnostics - Model 4</i>	55

1 Introduction

In May 2025, OpenAI announced its conversion from a non-profit entity to a public benefit corporation, opening the prospect of a future listing (Altman, 2025). Throughout its course, the firm has raised hundreds of billions of dollars in private capital, with subsequent funding rounds ranking among the largest private capital raises in history, while operating through a governance structure that concentrates control among a restricted set of partners, prohibiting both external nominations and ownership transfers (Mosco & Felicetti, 2024). If and when OpenAI enters public markets, the question becomes inevitable: will it adopt a dual-class share structure to preserve that control? Historical precedent suggests an affirmative answer. According to Ritter's IPO dataset (2026), in 2025 fully 41.1% of all U.S. IPOs adopted a dual-class structure, compared to approximately 7% in 2000. The OpenAI case illustrates the theoretical logic this thesis tests: a setting in which an expanding supply of late-stage private capital is associated with founders reaching the point of listing in a bargaining position consistent with the imposition of control structures that would have been unthinkable two decades ago; structures that, as they shift from exception to norm, may correspond to a structural reallocation of governance power from public investors to founding shareholders across the U.S. equity market.

Why do founders seek dual-class structures? For a firm whose long-term strategy may diverge from the expectations of quarterly-driven markets, dual-class structures offer a quasi-private governance arrangement while retaining access to the full depth of public capital. Dual-class structures allow founders to maintain control through shares carrying superior voting rights, while public investors receive shares with inferior or no voting power. The Ford family, for instance, retains over 40% of voting power while holding approximately 2% of equity (Gompers et al., 2010). This separation generates what the literature terms the “wedge” between voting power and cash-flow rights, quantified by Gompers et al. (2010) at an average of 60% voting against 40% economic ownership for insiders. Whether this divergence protects the founder's long-term vision (Goshen & Hamdani, 2016) or entrenches managerial control at the expense of outside investors (Bebchuk & Kastiel, 2017) remains contested; empirical evidence suggests that dual-class structures create value in young firms but erode it over time (Cremers et al., 2018; Kim & Michaely, 2019).

This diffusion did not occur in an institutional vacuum. Since the late 1990s, a series of regulatory, institutional and technological shifts expanded the U.S. private capital market substantially, with selected indicators rising several-fold: aggregate dry powder grew from approximately \$51 billion in 2000 to over \$280 billion by 2022, while aggregate late-stage deal volume, after contracting to a trough of approximately \$10 billion in 2003 in the aftermath of the dot-com bubble, expanded to a peak of approximately \$160 billion in 2021 before correcting toward the \$75-95 billion range by 2022-2023 (PitchBook Data & NVCA, 2017, 2024, 2026). The combination of deregulated private capital raising (Ewens & Farre-Mensa, 2020), crossover investor entry into pre-IPO rounds (Kwon et al., 2020), and a structural contraction of annual IPO volume (Gao et al., 2013; Doidge et al., 2017)

ensured that firms reaching public markets did so with credible outside options and a bargaining position shaped by the capital market conditions that preceded the listing.

Despite the centrality of this phenomenon, the relationship between late-stage private capital and governance choices at the moment of listing remains empirically underexplored. The existing literature predominantly focuses on firm-level determinants, such as founder characteristics, underwriter reputation, and industry affiliation, while largely overlooking the role of macro-financial conditions that may shape founders' bargaining power at the IPO stage. Aggarwal et al. (2022) represent the most relevant exception; however, their analysis spans the period 1994-2019, excluding the post-pandemic capital market cycle, and does not disentangle potential capital (dry powder) from the realized intensity of investment activity (deal volume). The present study addresses this gap by examining the association between late-stage deal volume and the adoption of dual-class structures in U.S. IPOs over the period 2000-2023. It interprets the association through a bargaining power channel, whereby an intensification of late-stage deal-making expands founders' outside options and strengthens their capacity to impose governance structures that preserve control at the moment of listing.

The research question is the following:

Is the adoption of dual-class structures in U.S. IPOs positively associated with the aggregate volume of late-stage private capital deal-making activity?

The study is explicitly associative and does not advance claims of causal identification.

To answer this question, a logistic regression model is estimated on a sample of 3,199 operating-company U.S. IPOs over the period 2000-2023, integrating Ritter's IPO database with WRDS/Compustat firm-level data and PitchBook-NVCA Venture Monitor private capital aggregates; the methodology is detailed in Chapter 3.

The study makes two contributions that the associative design supports. Empirically, it extends the temporal horizon of Aggarwal et al. (2022) to the period 2000-2023, incorporating the post-pandemic cycle, and it distinguishes between aggregate dry powder and realized deal volume as competing operationalizations of private capital availability, presenting empirical evidence consistent with a bargaining power channel that runs through active capital deployment rather than through the latent accumulation of uncommitted funds: when both measures are entered jointly, deal volume retains significance at the 5% level while dry powder attenuates to marginal significance, an empirical pattern that aligns the theoretical prior of the thesis with the precision of the estimates. Theoretically, it advances realized late-stage deal-making as a candidate upstream mechanism whose operation is consistent with the observed translation of private capital market conditions into governance outcomes at the moment of listing, a mechanism situated outside the domain of the public-market instruments built to discipline a separation of ownership and control whose structural form the moment of listing

now systematically inverts.

The thesis is structured as follows. Chapter 2 develops the literature review and formulates the research hypothesis. Chapter 3 describes the methodology. Chapter 4 presents the empirical findings. Chapter 5 discusses the results, limitations, and directions for future research. Chapter 6 concludes.

2 Literature Review, Theoretical Framework, and Hypothesis

This chapter examines the literature relevant to the research problem and develops the hypothesis for testing. The review is organized around three main ideas that together create a clear argumentative flow: from macro-financial conditions that may shape founders' bargaining position, to governance choices derived at the moment of listing, until the normative debate on the efficiency and legitimacy of such choices. Section 2.1 examines the structural change of the private capital market in the U.S. and how that affects decisions about listing. Section 2.2 analyzes empirical literature on the adoption of dual-class structures at the moment of IPO, emphasizing the adopters, the reasons, and how the market reacted. Section 2.3 discusses the theoretical debate between founder vision and agency costs, analyzing the life-cycle dynamics that determine whether dual-class structures create or destroy value. Section 2.4 identifies the research gap, delineates the conceptual framework, and formulates the research hypothesis.

2.1 The structural transformation of the U.S. private capital

The bargaining position that a founder carries into the IPO is shaped, before the listing begins, by the conditions of the private capital market that preceded it; since the late 1990s, a convergence of regulatory, technological, and institutional shifts has transformed those conditions in a direction theorized to reinforce founders relative to the counterparties they negotiate with at the moment of listing.

The genesis is regulatory. The National Securities Markets Improvement Act (NSMIA) of 1996 preempted state-level restrictions on private capital raising by qualified purchasers, igniting an exponential growth of venture capital and growth equity funds (Ewens & Farre-Mensa, 2020). Using fixed-effects regressions and instrumental-variable specifications, Ewens & Farre-Mensa demonstrate that variation in the supply of private capital at the sectoral level influences the probability that an IPO is founder-controlled, the size of the voting-cashflow wedge, and overall IPO incidence. Data from the PitchBook-NVCA Venture Monitor (PitchBook Data & NVCA, 2017, 2024, 2026) and the PitchBook Deals Screener (PitchBook Data, n.d.) confirm the magnitude of this transformation: aggregate dry powder, as the stock of capital committed to but not yet deployed by U.S. VC funds, rose from approximately \$51 billion at year-end 2000 to over \$280 billion by year-end 2022, following a near-monotonically increasing trajectory; late-stage deal volume, by contrast, as the flow of capital actively invested in Later Stage VC, Series C, and Series D transactions, exhibited a markedly cyclical profile, contracting to approximately \$10 billion at the post-bubble trough of 2003, recovering through the mid-2000s, expanding to a peak of approximately \$160 billion in 2021, and correcting toward the \$75-95 billion range by 2022-2023; deal counts, a complementary measure of deal-making intensity, increased fourfold across the same period. The central implication is that this expansion has the features of a structural shift rather than of a purely cyclical phenomenon: the cost of staying private appears to have diminished durably, altering the fundamental trade-off between listing and remaining

private; and it is precisely this altered trade-off that shapes founders' capacity to negotiate governance terms, including dual-class structures, at the moment of IPO.

Alongside this regulatory transformation, a technological shock reinforced the shift. The launch of Amazon Web Services in 2006 drastically reduced startup costs in exposed sectors (software, SaaS, online retail), producing a decline of approximately 20% in first-round investments and a near-doubling in the number of deals (Ewens et al., 2018). This shifted the VC investment model toward a high-volume, low-monitoring approach: more firms funded with smaller individual amounts, fewer board seats taken, and weaker investor oversight (Ewens et al., 2018). The result was a market with more capital competing for founders' attention and fewer constraints attached to it, a shift that expanded founders' bargaining power both during private rounds and at the moment of listing.

The effect of this shift was compounded by a third channel, opened by the entry of investors traditionally confined to public markets into the financing of late-stage private firms. Funds such as Fidelity and T. Rowe Price increasingly participate in pre-IPO rounds of high-growth firms, often alongside traditional venture capital funds; these crossover investors bring public-market capital into the private domain, and their presence is associated with higher private valuations, larger round sizes, and a reduced immediate need to go public (Kwon et al., 2020). The effect is to extend the pool of investors available to founders before the IPO: public-market investors increasingly enter the private financing domain before the IPO, rather than waiting for founders to reach the public market. This further deepens the outside options available and is associated with a shift in the bargaining equilibrium.

The aggregate consequence of these three dynamics is a structural contraction of the U.S. IPO market: the average number of annual listings fell from 310 in the 1980-2000 period to 99 in 2001-2012 (Gao et al., 2013), and the number of U.S. listed firms declined from 8,025 in 1996 to 4,102 by 2012 (Doidge et al., 2017). The firms that continue to list in this selective environment are theorized to do so from a position of strength, with concrete alternatives and a bargaining capacity inferred from the private capital they accessed before the IPO.

The literature converges on a shared theoretical reading: the transformation of the private capital market, through regulatory change (NSMIA), technological shock (cloud computing), and institutional shift (crossover investors), is argued to have shifted the power equilibrium toward founders. It is against this theoretical backdrop that the growth of dual-class structures at the moment of the IPO is examined in the section that follows.

2.2 The adoption of dual-class structures at the moment of the IPO

Having established the macro-financial conditions that reinforced founders' bargaining power, the analysis turns to the governance choices through which that power is exercised at the moment of listing. The literature documents who adopts dual-class structures, how they are engineered, and how

the market receives them; the magnitude of their growth is the empirical starting point.

The first large-scale analysis of American dual-class firms (1995-2002) reveals that approximately 6% of U.S. public companies, or about 8% of market capitalization, operated under a dual-class structure, with insiders holding, on average, 60% of voting rights but only 40% of cash-flow rights; a substantial wedge (Gompers et al., 2010). Firm value increases with insiders' cash-flow rights and decreases with their voting rights, suggesting that the separation between control and economic ownership generates measurable costs. This study established the pre-2000 baseline against which subsequent literature measures the evolution of the phenomenon. Both Gompers et al. (2010), whose sample spans 1995-2002, and Smart & Zutter (2003), whose sample spans 1990-1998 and is discussed below, cover periods that precede the structural transformation documented in Section 2.1, which makes the temporal extension of the present study to 2000-2023 particularly relevant.

The contribution that redefines the understanding of the subsequent phase is Aggarwal et al. (2022), who construct the first comprehensive database of dual-class IPOs in the United States (1994-2019, 5,907 IPOs), classifying each offering by share-class structure, type of controlling shareholder, and size of the wedge. The share of dual-class IPOs grew from approximately 5% to 28%, and nearly the entirety of this increase is attributable to founder-controlled firms, whose incidence on total IPOs rose from 3% to 19%. Founder wedges average 26.1 percentage points, significantly exceeding the 14.1 percentage points of other controllers; public shareholders, symmetrically, bear a negative wedge of -17.5 percentage points. The most relevant finding for the present study is the link Aggarwal et al. establish between governance choices and private capital market conditions. A one-standard-deviation increase in industry dry powder (approximately \$30 billion) is associated with a 3 percentage point increase in the probability of founder control; the 2006 cloud computing shock, which expanded the supply of private capital through reduced experimentation costs (Ewens et al., 2018), adds roughly 10 percentage points.

The consequences extend beyond the initial share-class choice. Greater founder bargaining power is associated with lower adoption of sunset provisions, clauses that provide for automatic conversion to a one-share-one-vote structure after a specified period or ownership threshold; Aggarwal et al. report reductions of 13.5 percentage points for time-based sunsets and 24 percentage points for ownership-threshold sunsets. The absence of sunset provisions matters because it transforms a governance choice made at the moment of listing into a potentially permanent allocation of control, amplifying the life-cycle concerns discussed in Section 2.3. Aggarwal et al. employ dry powder as their primary measure of private capital availability; the present study reads this choice as reflecting the data conventions of the literature rather than a developed theoretical argument about the channel, and proposes that dry powder, as the stock of potential investment, is theoretically distinct from the flow of actual deal-making activity through which founders' outside options are most plausibly constituted. Two markets with identical levels of dry powder may present fundamentally different bargaining environments depending on whether that capital is actively being deployed. This distinction, largely unexplored in the existing literature, motivates the present study's choice to employ

late-stage deal volume as its principal independent variable.

From the demand side, the way institutional investors perceive and discipline governance choices at the IPO stage reveals an asymmetry. While mature public firms have progressively reduced their use of classified boards and dual-class structures, IPOs move in the opposite direction, with controversial governance provisions increasingly concentrated among young, and often founder-led, issuers in the technology and biotechnology sectors (Field & Lowry, 2022). Institutional investors tolerate classified boards when information asymmetry is high but systematically oppose dual-class structures, especially when wedges are large (Field & Lowry, 2022). This opposition proves ineffective in practice: dual-class structures continue to proliferate despite investor resistance, consistent with the bargaining power channel documented by Aggarwal et al. (2022). The most visible institutional response came not from regulators or stock exchanges but from index providers: in 2017, S&P Dow Jones excluded new dual-class firms from the S&P Composite 1500 and FTSE Russell imposed a minimum voting rights threshold for index eligibility (Friedman et al., 2017). This exclusion, however consequential in its signaling, addresses a symptom rather than a cause, since index membership is determined post-listing and does nothing to alter the pre-IPO bargaining dynamics that shape the governance choice. That the market accommodates rather than penalizes these structures is consistent with this institutional pattern: dual-class IPOs carry an average underpricing of 11.9% against 13.7% for single-class peers, an approximately three-percentage-point gap that does not appear sufficient to deter adoption (Smart & Zutter, 2003).

The evidence reviewed in this section points to a market equilibrium in which founders, reinforced by the availability of private capital, impose governance structures that preserve their control without incurring pricing penalties sufficient to deter adoption, despite the opposition of institutional investors and index providers. The scale of this diffusion carries implications that extend beyond individual firm governance. If the trajectory documented by Ritter (2026) were to persist, a growing share, and potentially a majority, of newly listed firms would operate under structures that decouple economic ownership from voting control, a condition that Jackson (2018) characterized as a form of corporate royalty. For institutional investors, this translates into a growing share of the investable universe that offers exposure to economic risk without commensurate governance participation. The question is whether the public equity market can absorb this proliferation without eroding the accountability mechanisms on which market legitimacy has historically rested.

2.3 Founder vision and agency costs: the debate on efficiency

The literature does not supply a definitive answer, and the tension between the two dominant perspectives constitutes the core of the normative debate on dual-class governance, a debate that revisits the classical separation of ownership and control (Berle & Means, 1932) in an inverted form. Where the original problem was managers who control without owning, dual-class structures produce founders who both own and control, ostensibly resolving the classical agency problem. Whether this alignment genuinely serves all shareholders or merely relocates the agency cost, from the boardroom

to the public investor bearing economic risk without commensurate voice, is precisely what the literature contests.

The principal theoretical argument in favor of founder control draws on the concept of idiosyncratic vision, a specific mechanism within corporate governance theory formalized by Goshen and Hamdani (2016). In this framework, controllers pursue an entrepreneurial strategy they believe superior to market expectations; control is the mechanism that allows this vision to be implemented despite investor skepticism and information asymmetry, with benefits shared pro rata across all shareholders when the vision succeeds.

The strongest empirical footing for the idiosyncratic-vision defense rests on Jordan et al. (2016), whose propensity-score matching analysis of U.S. IPOs (1994-2011) documents that dual-class firms grow more rapidly than single-class peers (average annual revenue growth of 23% against 7.7%), maintain higher R&D intensity, and face markedly lower short-term market pressure; transient institutional ownership is 26.4% lower, takeover exposure is reduced, and analyst coverage is thinner. The strongest suggestive evidence sits in the subsample of firms that subsequently unify their share classes: revenue growth drops by 13 percentage points and R&D intensity by 1.5 percentage points once the dual-class structure is removed, while short-term pressure metrics rise, a within-firm pattern that is harder to reconcile with selection on adopter characteristics than with a protective function of the structure itself, although the design does not establish causal identification in a formal sense. The interaction with growth is asymmetric, since the Tobin's Q premium attaches only to high-growth dual-class firms and not to low-growth ones, which is the empirical qualifier the defense requires: the market rewards the combination of growth potential and governance protection, not governance protection standing alone.

The opposing perspective holds that perpetual dual-class structures are untenable (Bebchuk & Kastiel, 2017). The rationale is tied to a time horizon: benefits linked to the founder's protection are circumscribed to the phase in which the original vision is fresh and information asymmetry is at its maximum. Agency costs of misaligned control grow with the age of the firm, as the founder's initial advantages may erode, competitive conditions shift, and succession risk accumulates. Without a sunset mechanism, the founder remains entrenched even when their control is no longer efficient. Bebchuk and Kastiel argue, under a regulatory lens, that dual-class structures may be justified only if temporary and accompanied by effective sunset provisions.

This perspective is supported by empirical evidence consistent with a life-cycle hypothesis. Dual-class firms present a Tobin's Q 11% higher than comparable single-class firms at the time of IPO; however, this premium erodes within 6 to 9 years, turning into a discount of approximately 15% at maturity (Cremers et al., 2018). In parallel, the controllers' wedge grows from 16% one year post-IPO to 22% after five years and 26% after nine, a trajectory that signals increasing agency costs because the divergence between voting power and economic ownership, the condition under which private benefit extraction becomes possible, widens over time. Voluntary unifications of share classes are most

frequent between 3 and 5 years post-IPO, then decline. Kim & Michaely (2019) confirm this pattern with one of the most comprehensive datasets in the literature (920 firms, 8,455 firm-year observations, 1971-2015): young dual-class firms exhibit a valuation premium and operating performance comparable to single-class peers, but after approximately 12 years, labor productivity, margins, innovation, and Tobin's Q systematically decline. This reversal supports the life-cycle hypothesis and underscores the temporal dimension that distinguishes the optimistic and pessimistic readings of dual-class governance.

A relevant nuance emerges from recent evidence on the limits of formal control. Despite the voting power conferred by super-voting shares, approximately 12% of founder-CEOs of dual-class firms are removed, typically following a prolonged decline in Tobin's Q or a major legal risk event (Aran & Pollman, 2024). Even in the absence of formal sunset provisions, activist pressure, litigation risk, and board dynamics operate as implicit sunsets that limit, without eliminating, the control conferred by dual-class structures. For the present study, this evidence reinforces the idea that dual-class structures are not a permanent and immutable outcome, but a governance choice whose equilibrium depends, among other things, on the capital market conditions that favored their adoption in the first place.

The literature thus establishes that dual-class structures carry consequences that evolve over the life cycle of the firm; yet the governance architecture itself is determined at a single point, the moment of listing, where the capital market conditions documented in Section 2.1 exert their fullest influence.

2.4 Research gap, conceptual framework and development of the hypothesis

The existing literature provides a strong baseline and identifies a plausible theoretical channel, yet one that remains empirically incomplete. On one side, the structural transformation of the private capital market is documented (Ewens & Farre-Mensa, 2020; Ewens et al., 2018; Kwon et al., 2020), alongside the contraction of the IPO market (Gao et al., 2013; Doidge et al., 2017). On the other, the growth of founder-controlled dual-class structures and their relationship with private capital availability at the sectoral level is documented (Aggarwal et al., 2022). The empirical evidence, however, presents three specific limitations.

First, Aggarwal et al. (2022) cover the period 1994-2019, excluding the post-pandemic cycle of 2020-2021, which saw a sharp increase in both IPO volume and dual-class adoption, with the share reaching 33.1% in 2021 and 41.1% in 2025 (Ritter, 2026); this matters because the post-pandemic capital market environment, characterized by historically elevated levels of private capital deployment, represents the period in which the bargaining power channel would be expected to operate most intensely.

Second, the literature does not distinguish between potential capital (dry powder) and realized investment activity (late-stage deal volume), treating them as implicitly interchangeable when they

may operate through distinct mechanisms: dry powder reflects the stock of committed but undeployed capital, a series whose trajectory across the 2000-2023 period is near-monotonically increasing and therefore captures the secular accumulation of uncommitted funds rather than their cyclical deployment; deal volume, by contrast, captures the flow of active investment, a series whose year-to-year variation tracks the concrete financing alternatives actually available to founders in the period leading up to listing. The theoretical prior is that the bargaining power channel operates through the flow rather than through the stock: two markets with identical levels of dry powder can present fundamentally different bargaining environments depending on whether that capital is actively being deployed, and it is the rate of deployment, not the size of the reservoir, that founders observe and respond to in the negotiations that precede listing. This distinction is empirically testable, and its resolution carries practical relevance: if the channel operates through deal-making rather than uncommitted capital, then the appropriate measure of private-capital availability is the realized flow of late-stage transactions rather than the accumulated stock of committed but undeployed funds, and the empirical test that follows is set up to distinguish the two.

Third, the existing evidence operates at the industry-year level (Aggarwal et al., 2022), without testing whether the aggregate temporal dimension of late-stage capital availability, the overall market environment in which a founder approaches listing, independently shapes the probability of dual-class adoption at the firm level. The present study addresses all three: it extends the horizon through the post-pandemic cycle, it distinguishes dry powder from deal volume as competing operationalizations of private capital availability, and it tests the aggregate temporal channel that the industry-year literature does not isolate.

The conceptual framework of this thesis is based on the bargaining power channel identified by Aggarwal et al. (2022) and enriched by the literature on founders' outside options (Ewens & Farre-Mensa, 2020; Kwon et al., 2020). The mechanism operates as follows: an increase in late-stage deal-making activity expands the concrete alternatives available to founders, including additional private financing rounds, higher valuations, and the possibility of delaying the listing, thereby reinforcing their bargaining power.

This increased bargaining power is theorized to translate into a greater capacity to impose governance structures that preserve control, most directly through dual-class share structures. On the other side of this dynamic, the institutional counterweights that could in principle restrain founders' power appear structurally misaligned with the moment of the decision: the actors positioned to resist founder entrenchment either have diminished monitoring intensity before the listing (Ewens et al., 2018) or acquire opposition leverage only after it (Field & Lowry, 2022), while the governance architecture is crystallized at IPO itself. The efficiency debate (Bebchuk & Kastiel, 2017; Cremers et al., 2018; Kim & Michaely, 2019) tells us that the consequences of this choice unfold over the firm's life cycle, but the choice itself is made at the moment of IPO, precisely when founders' bargaining power is at its peak and the counterweights at their weakest.

Figure A1 provides a schematic representation of this channel, linking late-stage private capital availability to founder bargaining power and, ultimately, to the adoption of dual-class structures at the moment of IPO.

From this conceptual framework, the following research hypothesis is derived:

H1: A higher aggregate volume of late-stage venture capital deals is positively associated with the probability of dual-class structure adoption in U.S. IPOs.

This hypothesis predicts a positive association between the volume of late-stage deals at the aggregate annual level and the probability of dual-class adoption at the individual IPO level. The study does not advance causal claims; the objective is to identify an associative pattern consistent with the bargaining power channel developed in this chapter. Chapter 3 describes how the hypothesis was operationalized and tested.

3 Methodology

This chapter describes the empirical strategy adopted to test the hypothesis developed in Chapter 2. The discussion is organized as follows: Section 3.1 defines the research design. Section 3.2 specifies the unit of analysis. Section 3.3 details the sample construction process and data sources. Section 3.4 operationalizes all variables. Section 3.5 presents the model specification and the treatment of standard errors. Section 3.6 outlines the robustness and sensitivity analyses. Section 3.7 addresses validity, reliability, and limitations.

3.1 Research design

The present study adopts a quantitative, deductive research design, applied to structured secondary data organized as a stacked cross-section of IPO cohorts spanning 2000-2023, in which firm-level observations are nested within years and the principal regressor varies only across years. The approach is deductive: the hypothesis, derived in Chapter 2 from the conceptual channel of founder bargaining power, is tested through a logistic regression on a dataset of U.S. IPOs over the period 2000-2023. The study is explicitly associative and does not advance causal claims. What is investigated is whether a variation in the conditions of late-stage private capital, on a yearly aggregated basis, is associated with specific governance choices at the single firm level at the moment of listing. The alignment between research question, conceptual framework, hypothesis, and empirical strategy follows the logic of the design, even if the regressor is a coarse aggregate proxy for the channel it is meant to capture: the bargaining power channel predicts that greater late-stage deal-making activity strengthens founders' outside options, and the logistic model probes this prediction associatively by estimating the conditional probability of dual-class adoption as a function of aggregate deal volume, controlling for firm characteristics and industry composition. Bargaining power itself is not directly measured: late-stage deal volume operates as a macro-level proxy for the market conditions theorized to shape it, and the design therefore tests an association between observable inputs to the channel and an observable outcome, rather than the channel's internal operation.

3.2 Unit of analysis

The unit of analysis is the individual IPO, and each observation is a domestic company that successfully listed on a U.S. exchange between 2000 and 2023. The dependent variable, dual-class structure adoption, is measured at the individual firm level. The principal independent variable, late-stage deal volume, varies at the aggregate annual level, implying that all firms undergoing an IPO in a given year share the same value. This structure entails that the source of identifying variation is exclusively temporal, not cross-sectional within a given year; the model exploits differences in capital market conditions across the 24 years of the sample, not across industries within the same year. This is the core tradeoff of the design and the reason year fixed effects cannot be used: any year-level

absorption of variation removes the regressor entirely. The methodology is, in this respect, closer to a time-series identification with stacked cross-sections of IPO cohorts than to a conventional firm-level panel, and the inferential weight of the design rests on the 24 annual observations of the principal regressor rather than on the 3,199 firm-level observations of the outcome. The two econometric consequences are addressed in Section 3.5.1 and Section 3.5.2: first, the substitution of year fixed effects with a linear time trend to preserve the identifying variation; second, the clustering of standard errors at the year level to account for the within-year correlation induced by the aggregate regressor.

3.3 Sample construction and data sources

3.3.1 Population and time frame

The population of reference comprises all IPOs of operating companies completed on U.S. exchanges (NYSE, NASDAQ, or AMEX) during the period 2000-2023. The starting year is motivated by the structural transformation of the IPO market documented by Gao et al. (2013), who report a collapse in annual listing volume after 2000, from 310 (1980-2000) to 99 (2001-2012); this makes the study of firms that still choose to list in this selective environment particularly relevant. The ending year (2023) maximizes the contemporary relevance of the study and incorporates the post-pandemic cycle of 2020-2021, excluded from the existing literature (Aggarwal et al., 2022 stop at 2019).

3.3.2 Inclusion and exclusion criteria

All IPOs of operating companies listed on NYSE, NASDAQ, or AMEX are included. The following were excluded: SPACs (Special Purpose Acquisition Companies), REITs, closed-end funds, unit offerings, ADRs, best-efforts offerings, natural resource limited partnerships, and bank and S&L IPOs. These criteria replicate the standard in the IPO literature (Loughran & Ritter, 2004; Gao et al., 2013; Aggarwal et al., 2022). The exclusion process was applied to Ritter's full IPO dataset (originally compiled for Loughran and Ritter (2004) and subsequently updated annually by Jay R. Ritter at the University of Florida, with the most recent version dated 2026), which contains over 16,000 IPOs spanning 1975-2025. ADRs were identified and removed via the dedicated flag in the dataset. SPACs, acquisition vehicles, and shell companies were identified through a composite exclusion score that summed three independently scored criteria: name-pattern matching against a list of SPAC-typical strings (e.g., "Acquisition Corp," "Blank Check," "Merger Sub"), each match contributing one point; unit ticker suffix indicators (suffixes "U" or "WS"), each contributing one point; and an anomalous-age flag set to one when the firm's reported age at IPO was zero or negative, an empirical signature of acquisition vehicles. Observations with a cumulative score of two or higher were flagged for exclusion. The full list of name patterns and the scoring routine are documented in the supplementary replication file. Observations missing a CRSP PERMNO identifier were removed because PERMNO serves as the primary key for merging with WRDS/Compustat, from which industry classifications and firm-level financial variables are sourced (as detailed below); without a valid PERMNO, neither the SIC code nor the firm-size control can be constructed. The implications of this restriction for selection

bias are discussed in Section 3.7.5. The sample was then restricted to the period 2000-2023.

3.3.3 Data sources

The study integrates four main sources. The first is the IPO dataset of Jay Ritter (University of Florida, 2026), publicly accessible, which provides the dual-class classification (binary variable), listing dates, VC-backing status, founding year, and firm identifiers (ticker, CUSIP, CRSP PERMNO). The second is WRDS/Compustat (CRSP/Compustat Merged - Fundamentals Annual), accessed through institutional academic license, which provides the 4-digit SIC industry codes, total assets, net income, and sales for each firm matched via CRSP PERMNO. The third is PitchBook, employed through the Deals Screener filtered by deal type (Later Stage VC, Series C, Series D), geography (United States), and period (2000-2009), which provides the volume and count of late-stage deals for the years not covered by the Venture Monitor. The fourth is the PitchBook-NVCA Venture Monitor, which provides aggregate U.S. venture capital dry powder (capital overhang), sourced uniformly from the Q4 2025 edition and measured as of December 31 of each year, lagged to $t-1$ in this study to preserve the temporal ordering between the regressor and the outcome (since dry powder is a stock measured at the close of the fiscal year, whereas IPOs occur throughout the subsequent calendar year, using the contemporaneous end-of-year value would produce a regressor observed only after the full IPO cohort has already been realized), and the volume of late-stage deals for 2010-2023. Late-stage deal volume was sourced from multiple Venture Monitor editions depending on coverage: Q4 2016 (PitchBook Data & NVCA, 2017) for 2010-2012, Q4 2023 (PitchBook Data & NVCA, 2024) for 2013-2014, and Q4 2025 (PitchBook Data & NVCA, 2026) for 2015-2023; the Deals Screener served as the sole source for 2000-2009. Using multiple editions introduces a potential comparability issue, as PitchBook periodically revises historical deal classifications and valuations. Across overlapping years a discrepancy check, filed in the supplementary replication materials, was used to inform the selection rule: for years covered by Q4 2023 or Q4 2025 (2013 onward), the more recent edition is used; for 2010-2012, where the Q4 2016 Venture Monitor is the only available source, its values are retained. The magnitude of cross-edition divergence and its implications for inference are reported in Section 3.7.5.

3.3.4 Dataset integration

The integration of these sources required merging datasets operating at different levels of granularity. At the firm level, SIC codes and financial variables from WRDS/Compustat were matched to Ritter's dataset through the CRSP PERMNO identifier, which is present in both databases and provides a unique, time-invariant firm-level key. This matching process is deterministic and does not require name reconciliation, substantially reducing the matching bias that characterized the earlier PitchBook-based approach; residual coverage gaps remain where CRSP PERMNO assignment is incomplete or where the Compustat record is missing for the relevant fiscal year. The private capital variables (LateStageVol, DryPowder_t1, LateStageCount) were merged at the annual level, with each firm receiving the aggregate value corresponding to its year of listing.

3.3.5 Final sample

The sample attrition proceeds as follows. The Ritter dataset, restricted to the period 2000-2023, non-ADR, non-closed-end offerings with a valid CRSP PERMNO, yields 5,042 observations. After application of the exclusion scoring procedure (predominantly SPACs and acquisition vehicles), the sample reduces to 3,504. Missing values on the DualClass variable reduce it to 3,480; missing SIC codes from Compustat reduce it to 3,325; missing Age values reduce it to 3,250. A further 48 observations are lost to perfect separation in the logistic estimation, where 2-digit SIC groups with zero or universal dual-class adoption are automatically excluded because the likelihood function cannot be maximized for these groups; the affected sectors are those with very few sample firms (typically fewer than ten IPOs across the 24-year window), they collectively represent less than 1.5% of the analytical sample, and their exclusion is unlikely to alter the estimated association in a substantively meaningful way given that the bargaining-power channel is not theorized to operate differentially in those sectors. Three additional observations are lost to missing total assets data in Compustat, required for the FirmSize variable. The final analytical sample comprises 3,199 observations.

The cumulative reduction from 5,042 eligible IPOs to 3,199 thus reflects three substantively distinct sources: the exclusion scoring described in Section 3.3.2 (1,538 observations), missing values on essential firm-level variables in Compustat (230 observations: 24 DualClass, 155 SIC, 75 Age, 3 FirmSize), and 48 observations lost to perfect separation in the logistic estimation. The sample selection is therefore driven primarily by sample-frame definition (operating-company IPOs) rather than by data availability, and the residual missing-data attrition (7.1% of the post-exclusion sample) does not exhibit a temporal or sectoral pattern indicative of systematic non-randomness; the SPAC-and-acquisition-vehicle exclusion that drives the bulk of the upstream reduction systematically removes a class of issuer (blank-check shells) whose governance structures are categorically distinct from those of operating companies and whose inclusion would confound the dependent variable. The implications of this composition for selection bias are discussed in Section 3.7.5. Table A1 reports the complete attrition process.

3.3.6 Replicability and data access

Ritter's dataset is publicly available on the University of Florida website. WRDS/Compustat is accessible through institutional academic licenses. The Venture Monitor is published quarterly by PitchBook and NVCA and is available online. The filtering choices applied to the Deals Screener are documented in a supplementary replication file. Statistical analysis was conducted in Stata 18; all do-files and working datasets are available upon request.

3.4 Variable operationalization

3.4.1 Dependent variable

The dependent variable, *DualClass*, is a binary indicator equal to 1 if the firm adopts a dual-class structure at the moment of IPO, and 0 otherwise, sourced directly from the Ritter dataset (2026). In the pre-estimation sample ($N = 3,250$), 576 firms (17.72%) adopt a dual-class structure. This binary operationalization captures the extensive margin of the governance choice (whether a dual-class structure is adopted) but not the intensive margin (the size of the voting-cashflow wedge, or the presence and stringency of sunset provisions); the conceptual framework speaks to a broader dynamic of bargaining power and control retrenchment, of which adoption is the most observable manifestation but not its full extent, and the intensive-margin extension is identified as a direction for future research in Section 5.6.

3.4.2 Independent variable

The existing literature employs dry powder as the primary measure of private capital availability (Aggarwal et al., 2022). The present study departs from this convention and adopts the aggregate annual volume of late-stage deals (*LateStageVol*) as the principal independent variable, for theoretical reasons developed in Section 2.4: the bargaining power channel is theorized to operate through the concrete outside options available to founders, and what is hypothesized to reinforce a founder's position at the moment of listing is not the abstract existence of uncommitted capital on funds' balance sheets but whether that capital is actively flowing into late-stage deals, expanding the set of available financing alternatives in a manner that a founder can verify through observed valuations, round sizes, and the pace of competing offers. The variable is, accordingly, a macro-level proxy for a firm-level mechanism: it captures the market environment that founders observe, not the bargaining position any individual founder occupies, and the inferential leap from the macro proxy to the firm-level outcome is one of the design's structural limitations rather than one of its strengths. *LateStageVol* is measured as the aggregate annual dollar volume, in billions of dollars, of U.S. late-stage venture capital deals (Later Stage VC, Series C, Series D), sourced as documented in Section 3.3; across the 2000-2023 sample, the variable ranges from \$9.89 billion in 2003 to \$159.59 billion in 2021, with a mean of \$44.74 billion and a standard deviation of \$44.21 billion, reflecting the pronounced cyclicity that distinguishes deployed from committed capital and that constitutes the identifying variation the model exploits. The empirical basis for this choice, including a direct comparison of the two measures under the full specification, is presented in Chapter 4.

Dry powder (*DryPowder_t1*), measured as aggregate U.S. VC dry powder at December 31 of year $t-1$, in billions of dollars (mean \$98.44, SD \$50.90, range \$50.75 to \$281.27), is not included in the main model progression but is retained for a dedicated comparative analysis reported in Chapter 4. The one-year lag preserves temporal ordering: dry powder is a stock measured at fiscal year-end, IPOs occur throughout the subsequent calendar year, and the contemporaneous value would be observed only after the listing cohort had already been realized.

3.4.3 Control variables

Control variables are theory-driven and measured at the moment of listing. VCBacked is a binary indicator for pre-IPO venture capital funding (Ritter, 2026), capturing the presence of investors with documented governance preferences who may resist ceding control to founders (Field & Lowry, 2022). TechInd is a binary indicator for technology sector membership, defined using the 36 four-digit SIC codes identified by Loughran and Ritter (2004, updated), capturing the sectoral concentration of dual-class adoption documented by Aggarwal et al. (2022). Age is the number of years between the founding date and the IPO year (Ritter, 2026), serving as a proxy for firm maturity. FirmSize is the natural logarithm of total assets at the fiscal year-end preceding the IPO (WRDS/Compustat, Fundamentals Annual), capturing the scale-related variation in governance choices that firm age alone does not fully account for.

The combination of Age and FirmSize is an imperfect proxy for organizational maturity: Age captures elapsed time since founding without information on intervening pivots, restructurings, or changes in management, and FirmSize captures balance-sheet scale at a single point in time without information on growth trajectory or operational complexity; richer maturity proxies (revenue history, employee count, prior funding rounds) are not consistently available across the full sample period and are left for future work.

ROA, defined as net income divided by total assets (WRDS/Compustat), is included as a sensitivity control in a robustness specification (reported in Chapter 4) but not in the main model. The exclusion from the main specification reflects the position that profitability is not part of the bargaining-power channel; its inclusion as a robustness check responds to the alternative reading that profitability could nonetheless influence governance at the IPO stage through unrelated mechanisms (signaling firm quality to investors, conditioning underwriter terms, or shaping founder confidence at the bargaining table), and the test is whether the LateStageVol coefficient is sensitive to that channel even if the channel itself is conceptually distinct from the one the thesis tests.

A natural extension of the control set would include macro-financial conditions external to the venture capital market, in particular IPO-market temperature (annual aggregate IPO volume), equity-market returns (S&P 500 annual return), interest-rate level (10-year Treasury yield), and equity-market volatility (VIX). These variables co-move with both LateStageVol and the propensity of firms to list, and their exclusion from the main specification raises a legitimate omitted-variable concern: late-stage deal volume may be partially proxying for general market cycles rather than for the bargaining-power channel itself. The main specification's reliance on a linear time trend to absorb shared secular movement is a partial response to this concern but not a full one, since nonlinear macro fluctuations remain in the residual; the limitation is acknowledged transparently in Section 3.7.5, and the structural-robustness program in Section 3.6 (R6) probes the residual sensitivity through alternative trend functional forms in lieu of a fully specified macro-controls regression, which is identified as a direction for future research in Section 5.6.

Table A2 summarizes the definition, measurement scale, and data source for each variable in

the model, including both the dependent variable and the full set of controls.

3.5 Model specification and standard errors

3.5.1 Model specification

Given the binary nature of the dependent variable, the statistical method adopted is a logistic regression. The logit specification is preferred to a linear probability model on two grounds: it constrains predicted probabilities to the (0, 1) interval, which a linear probability model does not, and it accommodates the heteroskedasticity inherent in binary outcomes without requiring ex post correction. The linear probability model offers the offsetting advantage of directly interpretable coefficients, which the present design partially recovers by reporting average marginal effects in Section 4.3; a linear probability specification estimated on the same sample returns coefficients of consistent sign and significance and is reported as an additional sensitivity check in the robustness section.

The model is estimated through four specifications, each adding a component to the preceding one to isolate a progressively narrower source of identifying variation in the association between the independent variable and dual-class adoption.

Model 1 (baseline) examines the raw association without fixed effects:

$$P(\text{DualClass}_i = 1) = \Lambda(\beta_0 + \beta_1 \text{LateStageVol}_t + \beta_2 \text{VCBacked}_i + \beta_3 \text{TechInd}_i + \beta_4 \text{Age}_i + \varepsilon_i)$$

Model 2 adds industry fixed effects:

$$P(\text{DualClass}_i = 1) = \Lambda(\beta_0 + \beta_1 \text{LateStageVol}_t + \beta_2 \text{VCBacked}_i + \beta_3 \text{TechInd}_i + \beta_4 \text{Age}_i + \alpha_j + \varepsilon_i)$$

where α_j denotes fixed effects at the 2-digit SIC level. This specification accounts for cross-industry differences in dual-class adoption rates, such as the higher incidence documented in technology and biotechnology sectors (Aggarwal et al., 2022; Field & Lowry, 2022).

Model 3 adds a linear time trend:

$$P(\text{DualClass}_i = 1) = \Lambda(\beta_0 + \beta_1 \text{LateStageVol}_t + \beta_2 \text{VCBacked}_i + \beta_3 \text{TechInd}_i + \beta_4 \text{Age}_i + \beta_5 \text{Year}_t + \alpha_j + \varepsilon_i)$$

The inclusion of a linear time trend rather than year fixed effects is the central identification choice of the design and warrants explicit framing as such. Since *LateStageVol* varies exclusively at the aggregate annual level, year fixed effects would absorb the entirety of its variation, making estimation of β_1 impossible; the linear trend is the substitute that preserves the regressor at the cost of

imposing a functional form on the secular component of temporal variation. This is a substantive identification limitation: the linear trend absorbs only the linear secular component of co-movement between the regressor and the outcome, leaving any nonlinear time pattern in the residual. β_1 captures the association between within-trend fluctuations and the probability of dual-class adoption, net of industry composition and the linear trend, but it cannot be interpreted as net of nonlinear time-varying confounders (tech cycles, IPO waves, liquidity regimes, monetary-policy shifts) that co-move with deal volume; the linear trend therefore carries substantial inferential weight, and the robustness of the estimate to alternative trend specifications is examined in Section 4.6.

Model 4, the main specification, adds firm size:

$$P(\text{DualClass}_i = 1) = \Lambda(\beta_0 + \beta_1 \text{LateStageVol}_t + \beta_2 \text{VCBacked}_i + \beta_3 \text{TechInd}_i + \beta_4 \text{Age}_i + \beta_5 \text{FirmSize}_i + \beta_6 \text{Year}_t + \alpha_j + \varepsilon_i)$$

The four-specification progression should not be read as a sequence of increasingly clean identifications: the design remains associative throughout, the variation that identifies β_1 remains the within-trend year-to-year movement of an aggregate regressor against an aggregate outcome rate, and what the progression accomplishes is the absorption of progressively more confounding variation, not the establishment of a counterfactual.

The subscripts i , j , and t denote the firm, industry, and year respectively. Results are reported both as logistic coefficients and as average marginal effects, which express the change in the predicted probability of dual-class adoption per unit change in the regressor.

3.5.2 Standard errors and clustering

Because the principal independent variable (*LateStageVol*) varies at the aggregate annual level, all firms listing in the same year share the same value of the key regressor. This generates within-year correlation in the error terms that heteroskedasticity-robust standard errors alone do not address. To account for this, all specifications are estimated with standard errors clustered at the year level (24 clusters), following the recommendation of Cameron and Miller (2015) for designs in which the treatment varies at a level coarser than the unit of observation. Clustering at the year level inflates the standard errors of the aggregate-level regressor relative to plain robust standard errors, providing a more conservative basis for inference. The finite-cluster properties of this design, and the wild cluster bootstrap (Cameron et al., 2008) applied to address them, are discussed in Section 3.7.5 and the bootstrap results reported in Section 4.5.

3.6 Robustness and sensitivity analyses

To assess whether the main findings are sensitive to specific methodological choices, six robustness checks are conducted, complemented by a dedicated comparative analysis of dry powder

versus deal volume reported in Section 4.4 and by the wild cluster bootstrap of the main specification reported alongside the diagnostics in Section 4.5. Unless otherwise noted, all robustness specifications use clustered standard errors at the year level and include industry fixed effects, the linear time trend, and FirmSize alongside the relevant controls.

The first check (R1) replaces LateStageVol (dollar volume) with LateStageCount (number of late-stage transactions per year), testing whether the association holds when capital activity is measured by deal frequency rather than monetary volume. The second check (R2) re-estimates the main specification as a linear probability model rather than as a logit, providing directly interpretable coefficients on the probability scale and verifying that the central association is not an artifact of the logit's nonlinear link function. The third check (R3) adds ROA (net income divided by total assets) as an additional firm-level control, testing whether the inclusion of profitability alters the coefficient on LateStageVol. The fourth check (R4) restricts the sample to non-VC-backed firms ($N = 1,541$), testing whether the association between late-stage deal volume and dual-class adoption operates independently of the VC-backing channel, since VC-backed firms may face specific governance pressures that confound the relationship. The fifth check (R5) restricts the sample to non-technology firms ($N = 2,522$), testing whether the result is driven by the technology sector, where dual-class adoption is disproportionately concentrated.

The sixth set of checks (R6) probes the structural sensitivity of the main result to the functional form of the time trend, an issue identified in Section 3.5.1 as the central identification limitation of the design. Three alternative specifications are estimated: a quadratic trend (R6a), which allows curvature in the secular component of temporal evolution; a piecewise-linear trend with a break at 2010 (R6b), corresponding to the regulatory and institutional shifts that consolidate the post-NSMIA private capital market; and a sample restriction that excludes the boom years 2020-2021 (R6c), testing whether the central association survives when the most extreme observations of LateStageVol are removed. R6 is the diagnostically informative class of checks for this design: where R1 through R5 perturb measurement, functional form, and sample composition while holding the time-trend specification fixed (within-design robustness), R6 perturbs the time-trend specification itself (structural robustness against time-varying confounding). The distinction matters because within-design robustness alone cannot rule out the alternative reading that the central association is an artifact of the linear trend; only the second class of checks can.

Detailed results for all six robustness checks are reported in Section 4.6.

3.7 Validity, reliability, and limitations

The associative nature of the study imposes constraints on the strength of the inferences that can be drawn, which are acknowledged transparently.

3.7.1 Internal validity

Two threats are recognized: endogeneity and reverse causality. It is conceivable that industries with a higher propensity for dual-class adoption attract more late-stage capital, reversing the intended direction. Since the independent variable is measured at the aggregate annual level rather than at the industry level, this concern is partially mitigated by an argument from scale: the total volume of U.S. late-stage deals in a given year is determined by thousands of transactions across hundreds of funds, and the governance choices of the relatively small number of firms that reach the public market in that year are an insufficient driver of that aggregate. This is a plausibility argument that constrains the magnitude of any reverse-causal effect rather than a formal identification strategy that excludes it; reverse causality is not ruled out within the associative design, and the argument from scale should be read as partial mitigation rather than as an affirmative defense. More substantively, the design is exposed to a class of time-varying omitted variables for which industry fixed effects and a linear time trend provide insufficient absorption: aggregate IPO-market temperature, equity-market returns and volatility, monetary-policy regime, and investor sentiment all co-move with both private-capital deployment and the propensity of firms to list with founder-favorable governance, and any of them could partially or fully account for the observed association between LateStageVol and dual-class adoption. The robustness program in Section 3.6 includes alternative trend functional forms (R6) precisely to bound this threat, but bounding is not elimination; a fully specified macro-controls regression is not estimated within the present design and is identified as the most consequential extension for future research. The strongest form of the threat, in which the omitted variable is itself nonlinear in time and correlated with deal volume in ways a linear trend does not absorb, remains live within the present design. Lower-order threats (cultural norms within founder communities, geographic concentration of serial entrepreneurs, the influence of particular law firms advising on IPO governance) operate at firm and network level and are reduced but not eliminated by the controls and fixed effects in the main specification.

3.7.2 Construct validity

The proxies used to measure late-stage deal-making activity are imperfect. LateStageVol captures deals classified as Later Stage VC, Series C, and Series D in PitchBook, but does not encompass all forms of late-stage private financing; growth equity rounds outside these categories, or crossover investments structured outside traditional VC deal types, may be excluded. The aggregate annual measurement does not capture within-year variation or industry-specific dynamics. These limitations are inherent to the available data sources and are consistent with the proxies used in the literature (Aggarwal et al., 2022; Ewens & Farre-Mensa, 2020); their econometric implication is classical attenuation bias on the LateStageVol coefficient, since measurement error in a continuous regressor is well known to bias the estimated coefficient toward zero, which means the reported estimate should be read as a lower bound on the true association rather than as an unbiased point estimate. Beyond attenuation, the macro-level measurement also imposes a structural limit on the design's ability to identify the firm-level mechanism: even an unbiased estimate of the macro-to-firm-outcome association does not establish that the channel runs through the bargaining position of any individual founder, since multiple firm-level channels are consistent with the same aggregate pattern.

3.7.3 External validity

Results are generalizable exclusively to the U.S. IPO market during the period 2000-2023. Differences in regulatory regimes, notably the prohibition or restriction of dual-class structures in several European and Asian exchanges, limit extension to non-U.S. contexts without further investigation.

3.7.4 Reliability

Reliability in the present design has three distinct components, addressed in turn. The first is source replicability: Ritter's dataset is the standard reference in the IPO literature, the dual-class classification does not require hand-coding by the author, WRDS/Compustat data are deterministically matched via PERMNO (removing the subjective component of name-based matching while retaining the residual coverage limitations described in Section 3.3.4), and the PitchBook variables are documented with edition-level source attribution, including a discrepancy check sheet filed in the supplementary replication materials. The second is procedural replicability: all sample-construction operations, including the exclusion-scoring routine described in Section 3.3.2, are implemented in Stata 18 do-files filed as supplementary materials; the routine is deterministic conditional on the input dataset, and an independent re-execution should reproduce the analytical sample of 3,199 observations exactly. The third is inferential replicability: the wild cluster bootstrap reported in Section 4.5 uses a fixed random seed (documented in the do-file), so the bootstrap p-value is reproducible across runs without random fluctuation. The combination of public source data, deterministic sample construction, and seeded inference is what defines reliability in this design; what it does not defend against is the dependence of the central inference on the proprietary PitchBook editions cited in Section 3.3.3, whose underlying data are accessible only through institutional license and whose historical revisions across editions introduce a measurement-inconsistency component documented in the discrepancy check.

3.7.5 Scope conditions and limitations

The study operates with 24 year-level clusters, which is below the conventional threshold of 30 to 50 clusters recommended for reliable inference under asymptotic approximations (Cameron & Miller, 2015). To address this directly, the main specification is further subjected to a wild cluster bootstrap (Cameron et al., 2008) appropriate to the finite-cluster regime, the result of which is reported in Section 4.5. A complementary concern, given that the 24-year window includes years of unusually elevated deal-making activity (notably 2020-2021), is whether the central association is driven by a small number of boom years; this is addressed through a leave-one-year-out re-estimation of the main specification, in which the model is fitted 24 times with one calendar year omitted on each iteration and the distribution of resulting β_1 estimates is examined for stability. The procedure is reported in Section 4.6.

The bootstrap is a finite-sample correction to inference, not a remedy for the structural

limitations of the design (the macro-to-firm aggregation mismatch in Section 3.4.2, the linear-trend functional form in Section 3.5.1, and the cross-edition measurement inconsistency discussed below). Its inferential value should not be read as compensating for limitations of a different econometric class. The most consequential of the design's identification limitations is the dependence of the central inference on the linear functional form imposed on the time trend, which is treated as a substantive question rather than a technical one and probed directly in the structural-robustness program (R6) reported in Section 4.6; the complementary extension to a fully specified macro-controls regression is identified as a direction for future research.

A further scope condition concerns the measurement consistency of the principal regressor. The construction of LateStageVol from multiple PitchBook-NVCA Venture Monitor editions, documented in Section 3.3.3, introduces a measurement inconsistency whose magnitude varies across the sample period. For the years covered by both the Q4 2023 and Q4 2025 Venture Monitors (2015-2023), cross-edition divergence is small (maximum 5.99%, mean 2.68%); for the years covered by the older Q4 2016 Venture Monitor and subsequently revised (2013-2016), divergence is substantially larger (21-35%), reflecting PitchBook's reclassification of certain transactions between the late-stage and venture-growth categories. The selection rule documented in Section 3.3.3 (most recent edition prioritized for overlapping years; Q4 2016 Venture Monitor used only for 2010-2012, where it is the sole available source) addresses the asymmetry but does not eliminate it, and the residual measurement error represents a structural limitation of the regressor.

The associative nature of the analysis constrains the policy implications that can be drawn; any recommendation is formulated as a conditional observation rather than a causal prescription. The empirical results of the estimation strategy described in this chapter are presented in Chapter 4.

4 Empirical Results

This chapter reports the empirical results of the model specified in Chapter 3. Section 4.1 establishes the descriptive landscape of the sample and the bivariate structure of the key variables. Section 4.2 builds the central finding through four specifications that progressively absorb sectoral heterogeneity, secular growth, and firm-level scale. Section 4.3 translates that finding into economic magnitude and examines its interaction with venture capital status. Section 4.4 compares realized late-stage deal volume against aggregate dry powder as competing operationalizations of private capital availability. Section 4.5 reports the diagnostic checks that bear on the reliability of the estimates, and Section 4.6 the robustness of the finding across alternative specifications and subsamples.

4.1 Descriptive statistics and group comparisons

Table A3 reports descriptive statistics for the pre-estimation sample of 3,250 operating-company IPOs. The overall rate of dual-class adoption over the 2000-2023 period is 17.72%, rising from 6.39% at the start of the sample to 39.44% at its peak in 2022, with post-pandemic years remaining above one third. LateStageVol has a mean of \$44.74 billion (SD = \$44.21) and a temporal profile characterized by pronounced cyclicalities, ranging across the sample period from the post-bubble trough to the post-pandemic peak; DryPowder_t1 has a mean of \$98.44 billion (SD = \$50.90) and follows a near-monotonically rising trajectory. The contrast in the temporal shape of the two private-capital series, cyclical against secular, is the first descriptive fact that will acquire interpretive weight in the specifications below. Half the sample is VC-backed, one fifth belongs to the technology sector, and mean firm size is 5.53 in logarithmic scale.

Firms that adopted a dual-class structure at the moment of IPO differ from single-class issuers along the dimensions the model treats as central. They listed in years of substantially higher aggregate late-stage deal volume (\$66.1 billion versus \$40.1 billion) and higher lagged dry powder (\$122.9 billion versus \$93.2 billion); they were less frequently VC-backed (39.2% versus 53.1%); and they were larger at listing (mean log-assets 6.14 versus 5.40). These are raw differences, not evidence of the hypothesized channel, but the direction of each is descriptively aligned with it: the IPO cohorts most likely to adopt dual-class structures are those that reached public markets when private capital was most actively flowing, and when the founders reaching that moment were doing so from firms of above-median scale. The remaining group-level contrasts reported in Table A3 track in the same direction without altering the descriptive picture these central differences already establish.

The correlation matrix surfaces the structural feature of the data that defines the central identification constraint of the design. Both private-capital measures correlate strongly with the calendar year ($r = 0.730$ for LateStageVol, $r = 0.758$ for DryPowder_t1) and with each other ($r = 0.770$), the dependent variable correlates with the calendar year as well, and the joint upward trajectory of the regressor and the outcome across the 24-year window is the dominant feature of the bivariate landscape. This matters more than the orderly numbers suggest. Most of the variation in LateStageVol

is secular drift, not cyclical fluctuation, and the identification of β_1 in the specifications that follow depends entirely on the residual year-to-year variation that remains after the trend is absorbed; the risk of spurious correlation between two trending series is, in consequence, elevated, and the inferential burden carried by the linear time trend in Models 3 and 4 is correspondingly substantial. The specifications below should therefore be read with this constraint in mind: what the model identifies is not the gross co-movement of late-stage capital and dual-class adoption (which is large and largely shared with everything else trending upward over the period) but the residual association between their cyclical fluctuations around a linear path, and the economic content of that residual is what subsequent sections must test.

4.2 Main results: four specifications

Table A4 reports logistic regression results across four specifications, each estimated with standard errors clustered at the year level (24 clusters). Model 1 includes only the firm-level controls and the principal regressor. Model 2 adds industry fixed effects at the 2-digit SIC level. Model 3 adds a linear time trend. Model 4, the main specification, adds the logarithm of total assets. The progression is deliberate: each specification isolates a progressively narrower source of identifying variation, and the coefficient on LateStageVol is tracked across the sequence as successive sources of confounding variation are absorbed.

In Model 1 the coefficient on LateStageVol is positive and significant at the 1% level, reproducing the raw bivariate association between late-stage deal-making activity and the probability of dual-class adoption. Model 2 adds industry fixed effects, under which the coefficient rises rather than falls. The direction of the change is consistent with the sectoral distribution of dual-class adoption interacting with the sectoral distribution of late-stage capital flows in a way that suppressed the pooled estimate, but the data do not permit a stronger structural interpretation than this; the within-industry coefficient is what the specification identifies, and its larger magnitude relative to Model 1 should be reported as a fact about the estimate rather than as a finding about cross-industry heterogeneity. Within this specification TechInd turns negative, but the coefficient is not interpreted independently: the Loughran-Ritter technology indicator is defined at the 4-digit SIC level and sits inside 2-digit SIC fixed effects that subsume most of its variation, so what remains is residual within-2-digit variation that admits no clean substantive reading. The same overlap caveat applies to TechInd in all subsequent specifications.

Model 3 adds the linear time trend. The coefficient on LateStageVol attenuates from 0.0135 to 0.0043, which is the expected outcome given the joint upward trajectory of the regressor and the outcome and given that the trend now absorbs the linear component of that co-movement. What matters for the present chapter is not the attenuation itself but what the post-attenuation estimate is identifying on: the residual year-to-year fluctuation of late-stage deal volume around a linear path, and the corresponding residual year-to-year fluctuation of the annual dual-class adoption rate. The coefficient retains significance at the 1% level under year-clustered standard errors, but the inferential

weight now sits on a narrower band of variation than the unconditional series, and the substantive interpretation of the residual association is reserved for Chapter 5.

Model 4 adds firm size. The coefficient on LateStageVol moves from 0.0043 to 0.0038 (SE = 0.0008, $p < 0.001$ under year-clustered standard errors); firm size itself enters positively and significantly, confirming that the size composition of IPO cohorts is independently associated with dual-class adoption. The Model 3 to Model 4 transition does not eliminate the LateStageVol coefficient, but the stability of an estimate identified on a residual band of variation should not be over-read: the residual variation on which the model identifies remains limited (a 24-year aggregate series with strong shared trend with the outcome) and remains potentially confounded by time-varying factors not absorbed by the linear trend or by firm size, and the structural-robustness program in Section 4.6 is what bounds how seriously the estimate can be taken outside the assumptions of the present specification. The finite-cluster properties of this design and the wild cluster bootstrap that addresses them are reported in Section 4.5.

Among the controls, VC-backed firms are less likely to adopt dual-class structures, firm size and the linear time trend are positively associated with adoption, and firm age is not statistically distinguishable from zero. TechInd enters with a negative and statistically significant coefficient but is not interpreted independently for the overlap reason stated above.

4.3 Economic magnitude and marginal effects

Table A5 reports the average marginal effects computed from Model 4. A one-standard-deviation movement in LateStageVol, approximately \$44 billion, raises the predicted probability of dual-class adoption from 12.5% to 14.5%, an increment of approximately two percentage points holding all other variables at their means. The magnitude is modest in absolute terms and small relative to the year-to-year movement in the raw adoption rate, which oscillates by considerably more than two percentage points across the sample period; the residual association the model identifies should be read accordingly, as a marginal contribution within the broader temporal dynamic rather than as the dominant driver of cohort-level variation in dual-class incidence.

A conditional-margins analysis disaggregates the effect by venture capital status. The average marginal effect of LateStageVol is positive and statistically significant in both the non-VC-backed (0.000499) and VC-backed (0.000398) subsamples, both expressed per \$1B and equivalent to roughly 2.2 and 1.8 percentage points per one-standard-deviation movement respectively, with the non-VC-backed estimate slightly larger in magnitude. The association is therefore present in both populations and its direction is consistent with the bargaining-power channel; the substantive interpretation of the differential, including the question of whether the larger non-VC-backed coefficient reflects the absence of a counterweight to founder bargaining or a separate selection mechanism, is reserved for Chapter 5.

4.4 Dry powder versus deal volume

Table A6 compares DryPowder_t1 and LateStageVol as alternative operationalizations of private capital availability. Each variable is entered separately into the full specification, and then both are entered jointly. Included individually, the two measures return coefficients of similar magnitude and explain essentially the same share of variation in dual-class adoption, but they do so with markedly different statistical precision: LateStageVol is more precisely estimated than DryPowder_t1 under year-level clustering. When both are entered jointly, both coefficients attenuate, as expected from a pairwise correlation of 0.77; LateStageVol retains significance at the 5% level while DryPowder_t1 falls to marginal significance.

In comparative performance terms, the two measures yield estimates of similar magnitude when entered separately into the full specification, but the precision and the joint-specification behavior favor LateStageVol: it is more tightly estimated under year-clustered inference, and it retains 5% significance when both measures compete for the same residual variation while DryPowder_t1 attenuates to 10%-level significance. LateStageVol is therefore used as the principal measure throughout the remainder of the empirical analysis; the interpretation of why the flow measure outperforms the stock measure, and what this implies for the operationalization of private-capital availability in subsequent work, is developed in Chapter 5.

4.5 Diagnostics

Two diagnostics bear on the inferential properties of the main specification. The first concerns model discrimination: Model 4 achieves an area under the ROC curve of 0.7955, within the range conventionally associated with acceptable discriminatory capacity for binary-outcome models of this kind. The second concerns small-cluster inference: given that the cluster count of 24 falls below the 30-to-50 threshold conventionally recommended for reliable asymptotic cluster-robust inference, a wild cluster bootstrap (Cameron et al., 2008) was applied to the main specification with 9,999 Rademacher replications and year-level bootstrap clustering; the bootstrap returns a p-value of 0.014 ($z = 2.09$) on LateStageVol, which confirms that the central finding retains statistical significance at the 5% level under the correction most appropriate to the finite-sample regime, and which locates the honest inferential range of the estimate between the asymptotic $p < 0.001$ and the bootstrap 0.014.

Both diagnostics are inferential rather than identifying. The ROC curve speaks to the in-sample classification performance of the fitted model, not to the credibility of the bargaining-power interpretation of β_1 ; the bootstrap corrects the small-cluster properties of the standard error, not the structural weaknesses of the design (the linear-trend functional form discussed in Section 3.5.1, the macro-to-firm aggregation mismatch discussed in Section 3.4.2, and the cross-edition measurement inconsistency in the regressor discussed in Section 3.7.5). The diagnostics confirm that the inference is not an artifact of cluster asymptotics; they do not, and cannot, establish that the estimated association is identifying the channel the thesis posits. The structural questions the design cannot internally answer

are bounded, where they can be bounded at all, by the robustness program reported in Section 4.6.

4.6 Robustness and sensitivity

Tables A7a and A7b report the six robustness specifications described in Section 3.6 (with the structural-robustness check R6 estimated in three sub-variants, R6a through R6c), organized into a within-design block (R1 through R5) that perturbs measurement, functional form, and sample composition while holding the time-trend specification fixed, and a structural-robustness block (R6a through R6c) that perturbs the time-trend specification itself.

Within the first block, the coefficient on LateStageVol is positive, of comparable magnitude to the main specification, and statistically significant at conventional levels in every check. R1 replaces monetary volume with the count of late-stage transactions per year and recovers a coefficient of 0.000245 ($p < 0.01$); the association is therefore not an artifact of the monetary scale at which capital is measured. R2 re-estimates the main specification as a linear probability model (the probit specification originally proposed has been dropped from the program as redundant); the LPM coefficient on LateStageVol is consistent in sign and significance with the logit estimate, so the result does not depend on the distributional form of the link function. R3 adds ROA as an additional firm-level control and the coefficient on LateStageVol is unchanged (0.00386, $p < 0.01$); ROA itself enters with a negative sign that is not statistically distinguishable from zero, indicating that profitability is not independently associated with the governance choice in this sample. R4 restricts the sample to the 1,541 non-VC-backed firms and the coefficient remains positive and significant at the 5% level (0.00270); the association therefore operates in the subpopulation in which a confounding role of pre-IPO venture capital investors can be ruled out by construction. R5 restricts the sample to the 2,522 non-technology firms and the coefficient remains positive and significant at the 1% level (0.00296), of comparable magnitude to the main specification; the association is therefore not confined to the technology sector that has historically attracted the disproportionate share of empirical attention in the dual-class literature.

The second block is the diagnostically informative one. R6a replaces the linear trend with a quadratic specification; the squared term is dropped by the maximum-likelihood routine because year² is empirically collinear with year over the 24-year sample window, so the quadratic specification reduces, in practice, to the linear one. The reading of R6a is therefore not that the linear trend is the best of two functional forms tested but that the data do not contain enough variation to distinguish the two within the polynomial family. R6c excludes the 2020-2021 boom years from the sample and recovers a LateStageVol coefficient that strengthens to 0.00812 ($p < 0.01$), indicating that the central association is not driven by the post-pandemic outliers and that, if anything, those years moderate rather than inflate the pooled estimate. R6b is the check that pushes back. It introduces a piecewise-linear trend with a structural break at 2010, splitting the secular component of temporal evolution into a pre-2010 and a post-2010 segment; under this specification the coefficient on LateStageVol falls to 0.00175 (SE = 0.00109, $p = 0.110$) and loses statistical significance at conventional levels, while the

post-2010 trend term enters strongly ($\beta = 0.0824$, $p < 0.01$) and absorbs the share of cyclical variation that the linear trend leaves in the residual. The honest reading is that the central association in the main specification depends on a particular assumption about the functional form of the time trend, and that allowing the trend to bend at 2010 reassigns most of the explanatory load from LateStageVol to the post-2010 trajectory itself. This is the structural identification limitation flagged in Section 3.5.1, made empirically visible: the linear-trend specification is the working assumption of the design, not a tested feature, and R6b is the bound that defines what the design can and cannot speak to.

A complementary check addresses the temporal stability of the central estimate. The leave-one-year-out re-estimation of Model 4, motivated in Section 3.7.5, fits the main specification 24 times, omitting one calendar year from the sample on each iteration. The 24 resulting estimates of β_1 range from 0.00324 (when 2014 is dropped) to 0.00791 (when 2021 is dropped), with a mean of 0.00394 and a standard deviation of 0.00087; the coefficient is positive and of broadly comparable magnitude in every iteration, which indicates that the central estimate is not driven by any single calendar year in the sample. The widest excursion is to 0.00791 when 2021 is dropped, which is the mirror image of the R6c finding that excluding the 2020-2021 boom years strengthens rather than weakens the estimate.

The pattern of results across the six checks does not establish identification. R1 through R5 confirm that the central association is robust within the design: it does not depend on monetary scale (R1), link function (R2), profitability (R3), VC backing (R4), or the technology sector (R5). R6a and R6c confirm that the association is not an artifact of polynomial form within the linear family (R6a) and is not driven by the post-pandemic outliers (R6c). R6b reveals where the limit of the design lies: the central inference depends on the linearity of the time trend, and a piecewise alternative reassigns the explanatory load to the trend itself. These are the bounds the design provides; what they bound is the residual association on which the model identifies, not the underlying causal channel, and the substantive interpretation of the pattern is taken up in Chapter 5.

5 Discussion

This chapter interprets the findings reported in Chapter 4 against the conceptual framework developed in Chapter 2 and draws out their implications for the literature on private capital, IPO governance, and the institutional architecture of American public-market capitalism. Section 5.1 interprets the central association and evaluates the principal alternative readings that the associative design permits; Section 5.2 develops the study's methodological contribution, the distinction between dry powder and deal volume as competing operationalizations of private capital availability; Section 5.3 interprets the pattern of the robustness analyses, distinguishing the within-design block (R1-R5) from the structural-robustness block (R6); Section 5.4 articulates the normative implications of the evidence, organized as a sustained argument about the inversion of the Berle-Means paradigm and its consequences for the investable universe, the temporal structure of the governance decision, and the tipping point at which dual-class adoption ceases to be an exception and becomes a structural default; Section 5.5 discusses the limitations of the analysis; Section 5.6 identifies the most productive directions for future research.

5.1 Interpretation of the main result

The evidence reported in Chapter 4 supports Hypothesis H1 as formulated in Section 2.4: the aggregate annual volume of realized late-stage venture capital deals in the United States is positively and significantly associated with the probability of dual-class structure adoption at the moment of IPO, net of sectoral composition, the secular growth in the adoption rate, and firm-level scale. The economically interpretable magnitude, a rise of roughly two percentage points in the predicted probability of adoption per one-standard-deviation movement in LateStageVol, falls within the range previously documented in the empirical literature; Aggarwal et al. (2022) report approximately three percentage points per industry-year dry powder standard deviation, and the present study reports approximately two percentage points per aggregate-annual deal volume standard deviation; the point estimates differ in magnitude, with the gap plausibly attributable to the move from industry-year to aggregate-annual identification, since the latter cannot exploit cross-sectoral variation in capital-flow intensity. That the two estimates fall within the same broad order of magnitude despite substantial differences in measure (deal volume against dry powder), period (2000-2023 against 1994-2019), and level of analysis (aggregate annual against industry-year) is a form of conceptual alignment with prior evidence rather than a strict external replication; the present study locates itself within a small body of work that treats private capital market conditions as a relevant determinant of governance outcomes at the moment of listing, without claiming the cross-design convergence carries the inferential weight that direct replication would.

The interpretive question the finding raises is what the association is evidence of, given that the within-design robustness program (R1 through R5) returns the same sign, order of magnitude, and statistical significance across five distinct classes of perturbation, with the structural-robustness program (R6) reporting the boundary of the linear-trend assumption rather than additional

confirmation. The association is robust within this dataset and within the linear-trend specification family, which is a different statement than asserting the association would survive a fully different identification strategy; what follows interprets the within-specification robustness as evidence consistent with the bargaining-power channel rather than as evidence that the channel has been established.

The bargaining power channel developed in Chapter 2, drawing on Aggarwal et al. (2022), Ewens and Farre-Mensa (2020), and Kwon et al. (2020), posits a specific pathway: cyclical intensification of late-stage deal-making activity expands the concrete outside options a founder can verify in the period preceding the listing, including additional private financing rounds, higher valuations, and the possibility of delaying the listing altogether, and this expansion shifts the balance of power in the negotiation between founder and pre-IPO investors in a manner that increases the probability the founder can impose governance structures preserving control. The evidence assembled in Section 4.2 is consistent with the channel at the layer of the progression that carries the theoretical load: once the linear time trend in Model 3 absorbs the secular component of the co-movement between private capital conditions and dual-class adoption, what remains for identification is the year-to-year cyclical fluctuation of late-stage deal volume around its secular path, which is precisely the source of variation through which the channel is theoretically predicted to operate, not a by-product of the broader accumulation of private capital that Chapter 2 documents.

Before that interpretation can be accepted, three alternative readings of the evidence must be considered and evaluated on their own terms.

The first is reverse causality, the concern that expectations of a forthcoming wave of dual-class IPOs, anticipated by sophisticated private-market participants, generate the late-stage deal-making activity rather than being generated by it. The mechanism is not implausible in the abstract, but it is implausible at the scale at which LateStageVol is measured; the aggregate annual volume of U.S. late-stage deals is determined by thousands of transactions across hundreds of funds, and the governance choices of the relatively small number of firms that reach the public market in any given year constitute an insufficient driver of that aggregate. The direction of the causal arrow posited by the theoretical channel is, as a structural matter, the more plausible of the two.

The second alternative is compositional selection, the concern that late-stage deal volume shifts the distribution of firms reaching IPO toward larger, more mature, or more founder-controlled companies and that the governance association is therefore an artifact of this compositional shift rather than the signature of a bargaining channel operating on comparable firms. The robustness evidence bears directly on this alternative: R3 adds profitability as an additional firm-level control without altering the coefficient, the Model 4 inclusion of firm size leaves the LateStageVol coefficient nearly unchanged, and R5 returns an effect of comparable magnitude in the non-technology subsample, a population in which compositional drift toward founder-controlled firms is less pronounced than in the technology sector. The compositional reading is not formally excluded, but neither is it well supported.

The third alternative is founder-preference endogeneity, the concern that founders with strong preferences for control time their listings to high-deal-volume years, so that LateStageVol serves as a proxy for founder selection rather than for a bargaining channel operating on a homogeneous population; this reading would predict that the effect should weaken or disappear in the non-VC-backed subsample, where the institutional pressures that constrain founders are at their weakest and founder timing is least constrained, and R4 returns a coefficient that remains positive and significant in precisely this subsample, which runs against the prediction. None of these alternatives is formally excluded by an associative design, but the honest reading of the full pattern of evidence is that the bargaining power channel is the most parsimonious account consistent with the within-design robustness taken together, not that it is the only account the data permit.

The negative coefficient on VCBacked is consistent with the literature on pre-IPO governance pressure exerted by venture capital investors (Field & Lowry, 2022); the conditional-margins analysis in Chapter 4, which finds the LateStageVol effect positive and significant in both the VC-backed and non-VC-backed subsamples, indicates that the bargaining-power association does not depend on the presence or absence of venture capital backing as a sufficient condition. Whether the differential between the two subgroups reflects a constraint mechanism, a selection mechanism, or a combination of both is a question this design cannot resolve.

5.2 Deal volume and dry powder: measurement as methodological contribution

The direct comparison reported in Section 4.4, between lagged dry powder and realized late-stage deal volume as alternative operationalizations of private capital availability, is the study's principal methodological contribution to the literature. Aggarwal et al. (2022), whose empirical approach this study most closely builds on, employed dry powder as their primary measure without problematizing the choice between stock and flow; the empirical comparison assembled in Chapter 4 establishes that this distinction is not a matter of conceptual refinement but a measurable determinant of inference. When entered separately into the full specification, the two measures return coefficients of similar magnitude but markedly different precision, with LateStageVol estimated substantially more tightly under year-level clustering; when entered jointly, both attenuate as their shared variation is partialled out, but LateStageVol retains significance at the 5% level while dry powder falls to marginal significance, indicating that the portion of the variation uniquely attributable to realized deal-making dominates the portion uniquely attributable to the latent stock of uncommitted capital.

The theoretical reading of this empirical pattern is that the bargaining power channel transits through a flow, not through a stock. What expands the set of concrete alternatives available to a founder in the period preceding the listing is not the aggregate quantity of capital nominally available to private markets, a quantity that accumulates steadily on the balance sheets of funds regardless of the pace at which they deploy it, but the visibility of that capital in the form of deals actually being done, where visibility is constituted by the round sizes actually being raised, the valuations actually being

struck, and the pace of term sheets actually circulating in the period a founder is negotiating the terms of the listing. The cyclical structure of LateStageVol captures this flow; the secular structure of dry powder does not. The convergence of theoretical and statistical preference on the same measure is not a trivial methodological observation; it identifies a specific direction in which subsequent empirical work on private capital and IPO governance should move, reserving dry powder for the distinct conceptual role of capturing supply-side accumulation and adopting deal volume where the research question concerns the demand-side intensity of deployment. The recommendation rests on one joint specification, and its extension to sub-periods in which the stock-to-flow ratio differs substantially (pre-2010 against post-2010 being the natural division) is work for subsequent research; what the present evidence establishes is that the stock-flow distinction is a measurable, not a conceptual, determinant of inference in empirical designs that operationalize the bargaining power channel.

5.3 Robustness: the structure of the evidence

The robustness program reported in Section 4.6 varies in diagnostic weight across its two blocks. Within the within-design block, R1 and R2 test the result against measurement scale and functional form, perturbations that are standard hygiene rather than independent corroborations; R3 adds a profitability control whose null coefficient was the theoretical prior, since there is no plausible mechanism through which annual variation in aggregate late-stage deal volume would be confounded by IPO-cohort profitability. The diagnostic weight of this first block concentrates in R4 and R5, the two subsample restrictions that isolate the most theoretically salient confounding channels the design faces; the structural-robustness block (R6) is treated separately in the closing paragraph of this section. In each specification the coefficient on LateStageVol remains positive and significant, and its magnitude remains stable across the perturbations; it is this stability, rather than the presence of significance in any single column, that constitutes the empirical content of the finding.

The non-technology subsample restriction is the most diagnostically informative of the within-design block. The technology sector has historically absorbed the disproportionate share of empirical and normative attention in the literature on dual-class adoption, and the most natural skeptical reading of any study documenting a relationship between private capital conditions and governance structures is that the relationship is confined to technology IPOs and extended to the broader sample only by aggregation. The R5 coefficient, comparable in magnitude to the main specification and significant at the 1% level in a subsample that excludes more than seven hundred technology firms, locates the bargaining power channel outside the sector in which the literature's priors would most readily place it; the channel operates across the full range of non-technology sectors in the U.S. IPO market over the 2000-2023 period, not as a sectoral peculiarity of the industry that has dominated the empirical conversation. The R4 coefficient complements this evidence by establishing that the association holds in the population of firms from which pre-IPO venture capital investors, the natural institutional counterweight to founder bargaining power in the Field & Lowry (2022) framework, are absent by construction. R4 and R5 together define the evidentiary floor of the main result: the association is not confined to the technology sector, and it is not confined to firms whose governance is shaped by the

discipline of venture capital investors.

The structural-robustness block, R6, occupies a different evidentiary position. R6c (boom-year exclusion) and R6a (quadratic trend reducing to linear by collinearity) confirm that the central association is not driven by the post-pandemic outliers and is not distinguishable from the linear specification within the polynomial family. R6b (piecewise-linear trend) is the check at which the design's identification claim reaches its limit: the LateStageVol coefficient loses significance when the time trend is allowed to bend at 2010, and the explanatory load reassigns to the post-2010 trajectory itself. The interpretation, as developed in Chapter 4, is not that R6b refutes the central finding but that it locates the boundary of what the present design can speak to; the within-design robustness summarized above retains its evidentiary value precisely because it operates within the linear-trend specification that R6b reveals to be a working assumption rather than a tested feature.

5.4 Normative and market implications

The findings of this thesis, read alongside the trajectory documented by Ritter (2026), open onto a set of questions about governance dynamics whose scope, if the conditional reading developed below is accepted, extends beyond the individual firm; the implications presented in the sections that follow are formulated as conditional readings of the evidence rather than as conclusions the associative design independently establishes, and each is anchored to the specific empirical pattern from which it draws. The sections that follow develop four implications in sequence: the inversion of the Berle-Means paradigm and its consequences for the instruments of corporate discipline; the temporal asymmetry between the moment of the governance decision and the horizon over which its consequences unfold; the recomposition of the investable universe and the structural limits of the corrective mechanisms institutional investors have so far deployed; and the tipping-point dynamics by which dual-class adoption, once an exception, is approaching a structural default.

5.4.1 The inversion of the Berle-Means paradigm

In 1932, Berle and Means diagnosed the fundamental problem of the modern corporation in the separation of ownership and control: dispersed and passive shareholders on one side, professional managers exercising effective control without holding a significant stake in the capital on the other. The regulatory edifice constructed in the nine decades that followed, from the Securities Exchange Act of 1934 to the Sarbanes-Oxley Act of 2002, from the fiduciary duties of the board of directors to the rise of institutional investor activism, was built to respond to that problem: to align the incentives of the manager with those of the shareholder, to bridge the informational gap, and to ensure that those who control the firm are accountable to those who bear its economic risk.

The proliferation of dual-class structures yields a configuration that, on the reading developed in the literature (Aggarwal et al., 2022; Jackson, 2018) and presupposed by the untenable-case argument of Bebchuk and Kastiel (2017), inverts the terms of the classical problem. The founder who

adopts a dual-class structure at the moment of IPO retains voting control even as the share of economic ownership held outside the founding group expands through the listing; in the limit, this configuration concentrates voting power in a holder whose economic stake may be a small fraction of total equity, and it leaves public investors bearing economic risk without commensurate governance power. The average wedge between voting rights and cash-flow rights, quantified by Gompers et al. (2010) at 60% voting against 40% economic ownership for insiders and documented by Aggarwal et al. (2022) at 26.1 percentage points for founders, has ceased to be a residual anomaly in a market where Ritter (2026) records 41.1% of 2025 IPOs adopting the structure. What the evidence of this thesis adds to that existing reading is the identification of a candidate upstream mechanism: if the findings of Chapter 4 are indicative of a persistent relationship between private capital market conditions and governance choices at the moment of listing, and if the factors that have fueled the expansion of the late-stage venture capital market, the post-NSMIA deregulation documented by Ewens & Farre-Mensa (2020), the reduction in experimentation costs documented by Ewens et al. (2018), and the entry of crossover investors documented by Kwon et al. (2020), are structural rather than cyclical, then the inversion observed in the present sample window may persist beyond it, with the bargaining architecture of new listings located upstream of the public-market instruments built to constrain governance choices at the moment of issuance. The 2000-2023 evidence cannot establish that this configuration is permanent or structural in the strong sense, only that the pattern observed within the sample is consistent with the conditional reading set out above.

The theoretical implication runs deeper than the descriptive observation. The regulatory and institutional infrastructure of the American equity market was constructed to solve the agency problem in its classical form; it was not designed to manage the inverse problem, the one in which the controller is also the owner and the agency cost falls entirely on the public investor who lacks a voice. The governance instruments that function within the Berle-Means paradigm, the shareholder vote, the independent board of directors, the proxy fight, the hostile tender offer, presuppose that voting power is proportional to economic ownership; once this proportionality is structurally severed, the disciplining function of those instruments is correspondingly weakened, and the literature on dual-class governance (Bebchuk & Kastiel, 2017; Jackson, 2018) characterizes this weakening as the analytical core of the contemporary critique. The board of directors remains formally in place, but the founder who holds a majority of votes with a minority of capital is, as a practical matter, insulated from shareholder-vote removal under most share-class architectures, although fiduciary-duty actions, board dynamics, and contractual sunset provisions where present can introduce limited corrective channels; the hostile tender offer is structurally constrained when the controller cannot be displaced through capital-market acquisition; institutional activism is reduced largely to reputational pressure, with limited formal leverage. Aran and Pollman (2024) document that approximately 12% of founder-CEOs of dual-class firms are removed, typically following a prolonged decline in Tobin's Q or a major legal risk event, but this percentage, while significant as a demonstration that control is not absolutely impenetrable, confirms that in 88% of cases the founder remains in place regardless of the will of public shareholders. The classical instruments of corporate discipline were calibrated for a world in which misalignment between ownership and control was a question of incentive design; the literature

on dual-class governance argues that they are less well-suited to a world in which the misalignment is a question of voting-rights architecture established once, at the moment of listing, and not meaningfully revisable thereafter, and the present evidence is consistent with that reading without being able to test it directly.

5.4.2 The temporal asymmetry of the governance choice

A second implication concerns the structural misalignment between the moment of the decision and the time horizon over which its consequences unfold. The dual-class governance structure is crystallized at a single point in time, the moment of IPO, which is also the point at which the founder's bargaining power is at its maximum and the institutional counterweights at their minimum. The institutional counterweights operate in the wrong temporal domain: pre-IPO investors have moved toward lighter monitoring (Ewens et al., 2018) precisely as founder bargaining power has strengthened, and post-IPO institutional opposition (Field & Lowry, 2022) operates after the governance architecture has already been set.

The consequences of this choice, however, do not exhaust themselves at the moment of listing; they unfold across the entire life cycle of the firm, and the empirical evidence on that life cycle points to material long-run governance costs. The initial valuation premium of dual-class firms, documented by Cremers et al. (2018) as a Tobin's Q approximately 11% higher than comparable single-class firms at the time of IPO, erodes within six to nine years and turns into a discount of approximately 15% at maturity; the controllers' wedge grows from 16% one year post-IPO to 22% after five years and 26% after nine, a trajectory that signals rising agency costs over the very horizon across which the initial valuation advantage decays; labor productivity, margins, innovation, and Tobin's Q systematically decline after approximately twelve years (Kim & Michaely, 2019). The empirical pattern documented in the cited life-cycle literature is that of a temporal asymmetry: the decision is taken at the moment of the founder's maximum bargaining strength, when the benefits of idiosyncratic vision are at their peak and agency costs at their minimum, but the most severe consequences surface decades later, when the original vision may have eroded, competitive conditions may have shifted, and succession risk has accumulated, yet the governance architecture remains unchanged.

Bebchuk and Kastiel (2017) argued, under a regulatory lens, that dual-class structures can be justified only if temporary and accompanied by effective sunset provisions; the evidence of Aggarwal et al. (2022) documents that greater founder bargaining power is associated with a 13.5 percentage point reduction in the adoption of time-based sunsets and a 24 percentage point reduction in the adoption of ownership-threshold sunsets. The bargaining power channel, in other words, does not only increase the probability that the dual-class structure is adopted; it also reduces the probability that the structure contains the corrective mechanism designed to attenuate its costs over time. The present thesis does not test the determinants of sunset provisions; the Aggarwal et al. evidence is reported here as prior literature whose pattern, if extended in subsequent research to the post-pandemic period covered by this study, would compound the temporal asymmetry described above by linking the

conditions that increase dual-class adoption to those that simultaneously reduce the adoption of the corrective mechanism designed to attenuate its long-run costs.

5.4.3 The recomposition of the investable universe and the limits of index exclusion

For institutional investors, the proliferation of dual-class structures among newly listed firms is associated, through index construction mechanics, with a progressive alteration in the composition of the investable universe; the implication of that alteration extends, on the reading developed below, beyond the governance of any individual firm. Pension funds, index funds, and exchange-traded funds construct their portfolios on indices that reflect the market capitalization of listed firms; to the extent that a growing share of these firms operates under structures that sever economic ownership from voting control, the institutional investor is progressively exposed to a universe in which governance participation is structurally limited. The shareholder vote, the cardinal instrument of participatory corporate governance, becomes formally available but substantively irrelevant when the founder holds super-voting shares sufficient to determine any resolution regardless of the will of the remaining shareholders.

The most visible institutional response has been the exclusion of new dual-class firms from indices. In 2017, S&P Dow Jones excluded new dual-class firms from the S&P Composite 1500 and FTSE Russell imposed minimum voting rights thresholds for index eligibility (Friedman et al., 2017). Index exclusion imposes an ex-post cost on adoption through institutional demand effects, but it operates downstream of the listing and does not alter the pre-IPO bargaining dynamics with which the present evidence associates the governance choice; its deterrent efficacy is, in consequence, structurally limited.

The regulatory paradox that emerges is one of domain asymmetry. The existing corrective instruments operate entirely within the domain of the public market, whether index exclusion, institutional investor opposition, or pressure for sunset clauses, while the cause that the findings of this thesis associate with the governance choice operates within the domain of the private market. So long as regulatory and institutional attention remains confined to the governance provisions of issuers at the moment of listing, it does not intercept the structural forces that feed founders' bargaining power in the pre-IPO phase; a realignment of corrective capacity with the upstream structure suggested by the present evidence would require an approach that treats the private capital market and the public market as components of an integrated institutional system rather than as separate compartments with distinct regulatory frameworks, an approach that is not yet visible in the policy landscape.

5.4.4 Dual-class adoption at the tipping point

The final implication concerns the tipping point at which the dual-class phenomenon may transition from a governance exception to a structural default. The literature has historically treated dual-class adoption as a deviation from the one-share-one-vote norm: Gompers et al. (2010) reported

approximately 6% of U.S. public firms in the 1995-2002 period, and Aggarwal et al. (2022) reported approximately 5% of IPOs in the 1990s. The 2025 adoption rate of 41.1% reported by Ritter (2026) lies outside the empirical sample of this thesis, which ends in 2023, and is cited here as contextual motivation rather than as evidence the present design tested. Read as context, the rate's proximity to the majority threshold marks a shift in the interpretive register through which the phenomenon can be understood: a governance provision adopted by 6% of issuers can plausibly be read as a peculiarity of a self-selected subpopulation, while one adopted by above 40% of issuers requires a different kind of explanation, a register the present thesis raises but does not independently establish. The cost of the choice, moreover, remains low; Smart and Zutter (2003) documented that dual-class IPOs over the 1990s carried average underpricing of 11.9% against 13.7% for single-class peers, a gap the market does not widen sufficiently to discourage adoption. A self-reinforcing dynamic, in which growing adoption rates progressively normalize the structure and lower the marginal cost of adoption for subsequent founders, is one possible interpretation of the trajectory but is not tested by the present design; whether such a feedback effect operates, and at what speed, is a question for subsequent research that the descriptive trajectory makes increasingly relevant.

In the most cited articulation of the regulatory critique, Jackson (2018), then a Commissioner of the Securities and Exchange Commission, characterized perpetual dual-class structures as a form of corporate royalty, an aristocracy of control in which power is transmitted not through market dynamics but through share structure, regardless of competence or performance; the present thesis cites this characterization to mark the analytical register at which the contemporary debate is conducted, not to endorse it as the descriptive verdict on dual-class adoption in the 2000-2023 sample. Whether the public equity market can absorb a proliferation of dual-class structures without eroding the accountability mechanisms on which its institutional legitimacy has historically rested, and whether the regulatory paradigm built for the Berle-Means problem remains adequate to a market in which the bargaining architecture has shifted upstream of public-market instruments, are questions the present evidence raises rather than answers. The associative design of this thesis cannot license policy conclusions on either; what the empirical test independently establishes is the within-design association reported in Chapter 4, and the questions raised in this section are advanced as conditional implications of that association read alongside the literature on post-NSMIA private capital expansion, rather than as findings the design itself supplies.

5.5 Limitations

The thesis has five main limitations. The most consequential is the absence of a direct measure of the theoretical mediator, founder bargaining power, on which the entire interpretation of the central finding depends; this limitation is developed in full as the fifth limitation below, and it is flagged here at the outset because it conditions the reading of the four limitations that precede it.

The first limitation is the 24-cluster problem. Standard errors are clustered at the year level because the principal independent variable varies at the aggregate annual level, which generates

within-year correlation in the error terms that heteroskedasticity-robust standard errors alone cannot address; the cluster count of 24 falls below the conventional threshold of 30 to 50 recommended for reliable asymptotic cluster-robust inference (Cameron & Miller, 2015), and the asymptotic p-values reported in Chapter 4 rest, therefore, on an approximation whose accuracy in the finite-sample regime cannot be taken for granted.

A wild cluster bootstrap following the procedure of Cameron et al. (2008), with 9,999 Rademacher replications and year-level bootstrap clustering, was applied to the main specification and is reported in Section 4.5; the bootstrap returns a p-value of 0.014 on LateStageVol, which is, as expected in small-cluster regimes, roughly an order of magnitude larger than the asymptotic $p < 0.001$. The honest inferential statement is therefore that the coefficient is significant at the 5% level under the inference regime most appropriate to the data, not at the 0.1% level under an asymptotic approximation known to be anti-conservative when the cluster count is small.

The 24-cluster regime remains the most structural constraint of the research design, and the bootstrap correction, while ensuring that this constraint does not invalidate the central inference, does imply that the scale of statistical confidence extractable from these data is bounded by the coarseness of the clustering structure; no post-estimation technique can fully substitute for the richer identifying variation that the industry-year disaggregation or instrumental-variable design discussed in Section 5.6 would provide.

The second limitation is the associative nature of the research design. The within-design robustness reduces but does not eliminate the threats posed by reverse causality, compositional selection, and founder-preference endogeneity discussed in Section 5.1; the distance between a documented association and a causal claim remains substantial, and the structural-robustness check R6b in Section 4.6 made empirically visible the bound that this distance imposes on the present specification.

The third limitation concerns the construct validity of LateStageVol. The variable captures deals classified as Later Stage VC, Series C, and Series D in PitchBook, but it does not encompass all forms of late-stage private financing; growth equity rounds outside these categories, crossover investments structured outside traditional venture capital deal types, and secondary transactions in pre-IPO firms are not captured. The aggregate annual measurement, moreover, does not capture within-year variation or industry-specific dynamics that could differentiate the effect of private capital on governance across industries with distinct financing profiles; if the bargaining power channel operates with varying intensity within the year or across sectors with distinct venture capital intensity, the annual aggregate measurement introduces measurement error that may attenuate the estimated coefficient toward zero under classical errors-in-variables conditions, although the direction of the bias is not guaranteed in nonlinear models or when measurement error is correlated with other regressors.

The fourth limitation concerns external validity. The results are generalizable exclusively to the

U.S. IPO market during the period 2000 through 2023. Differences in regulatory regimes, notably the prohibition or restriction of dual-class structures on several European and Asian exchanges, limit the extension of the findings to non-U.S. contexts; within the U.S. market, the sample excludes SPACs, REITs, closed-end funds, ADRs, and best-efforts offerings, so the results reflect the governance dynamics of operating companies accessing the market through traditional IPOs, not the governance dynamics of the universe of all new listings.

The fifth and final limitation is the absence of a direct measure of founder bargaining power. The theoretical channel connecting late-stage deal volume to dual-class adoption transits through bargaining power as a mediator, but bargaining power is neither directly observed nor measured in the model; LateStageVol is a proxy for the market conditions that reinforce bargaining power, and the documented association is consistent with the hypothesized channel but does not demonstrate its internal operation. A research design capable of directly measuring founder bargaining power, for instance through the terms of pre-IPO term sheets or the structure of preceding financing rounds, would permit a more direct test of the mechanism and is the most theoretically informative extension of the present work.

5.6 Directions for future research

The limitations identified in the preceding section structure the most productive directions for subsequent research, and the present section organizes them in ascending order of ambition, from the industry-year disaggregation that would tighten the identification of the present specification to the comparative-international design that would require new data infrastructure.

The first direction concerns the temporal and sectoral granularity of the independent variable. An analysis that disaggregates late-stage deal volume at the industry-year level would allow testing whether the bargaining power channel operates with varying intensity across industries and within the year, overcoming the limitation of the aggregate annual measurement imposed by the data structure of the present study; the availability of PitchBook data at the individual-transaction level makes this disaggregation feasible, and its implementation would permit the inclusion of industry fixed effects, year fixed effects, and industry-by-period fixed effects (with periods defined more coarsely than the year, for instance pre- and post-NSMIA, or pre- and post-2010) that together absorb cross-sectoral and intertemporal heterogeneity without absorbing the industry-year variation in the principal regressor itself, producing a more stringent test of the hypothesis than the present aggregate-annual specification supports.

The second direction concerns causal identification. The associative design could be strengthened through quasi-experimental strategies; Ewens and Farre-Mensa (2020) exploit the sectoral variation in the supply of private capital induced by the National Securities Markets Improvement Act as an instrumental variable for private-market capital supply, and an analogous approach could exploit exogenous shocks to late-stage deal-making volume, whether regulatory

changes affecting the ability of funds to complete deals in specific sectors or periods, or identifiable shifts in limited-partner commitment behavior, to isolate the exogenous component of LateStageVol and test the bargaining power channel under more stringent identification conditions than the present design supports.

The third direction concerns the direct measurement of founder bargaining power, which in the present design remains an unobserved theoretical mediator. Several specific operationalizations are tractable with deal-level data that is becoming progressively more accessible: the ratio of founder voting rights to founder cash-flow rights at each pre-IPO financing round (a within-firm trajectory that captures the evolution of bargaining position across rounds); the number of competing term sheets received in the round preceding the IPO (a direct count of outside options); the valuation step-up between the last private round and the IPO offer price (a market-based measure of how much of the founder's bargaining strength is converted into pricing); and the presence and stringency of investor protective clauses (board seats, redemption rights, anti-dilution provisions) in the latest term sheet, the inverse of which proxies founder leverage. Constructing any one of these from deal-level data would permit a direct test of the channel that the macro-level proxy of the present design can only approximate; the growing availability of such data, through the expansion of PitchBook coverage and analogous databases, renders this direction increasingly tractable for empirical research over the coming years.

The fourth direction extends the temporal focus to the post-pandemic cycle and its correction. The sample of the present study ends in 2023; the years 2024 through 2026 have seen a significant contraction in IPO volume and a correction in late-stage venture capital valuations, while the dual-class adoption rate has continued to rise (Ritter, 2026, reports 41.1% in 2025), and extending the sample to this period would permit testing whether the relationship between deal volume and governance holds under restrictive capital market conditions, or whether dual-class adoption has reached a level of diffusion such that it has become a structural default independent of the cyclical conditions of the private capital market.

The fifth and final direction concerns the comparative international extension. Differences in regulatory regimes governing dual-class structures across jurisdictions, whether the prohibition in certain European exchanges, the recent liberalization in others, or the restrictions in several Asian markets, offer an opportunity to exploit institutional variation as a source of identification; a comparative analysis that contrasts markets under different regimes but exposed to similar private capital flows could isolate the effect of regulatory conditions from that of market conditions, contributing to a more complete understanding of the factors that determine IPO governance at the global level.

These five directions form a coherent research program: industry-year disaggregation, quasi-experimental identification, direct measurement of the mediator, post-2023 extension, and international comparison. The present thesis establishes the associative foundation on which such a program can be

built.

6 Conclusion

This thesis asked whether the aggregate volume of realized late-stage venture capital deal-making activity in the United States is positively associated with the probability of dual-class structure adoption in U.S. IPOs over the period 2000 through 2023, and the evidence assembled supports that association as a systematic feature within the linear-trend specification family, with the boundary of that conditional support reported transparently in Section 4.6. Cyclical variation in late-stage deal volume, net of the secular growth of the phenomenon, of sectoral composition, and of firm-level scale, is positively and statistically significantly related to the probability that a founder adopts a dual-class structure at the moment of listing; the association is recoverable under alternative measurement of the independent variable, under alternative functional form, under the addition of a profitability control, and under restriction of the sample to the non-VC-backed and non-technology populations that constitute the most theoretically salient confounding-channel concerns the design faces. It is the stability of the effect across these perturbations, rather than the magnitude of its estimate in any single column, that constitutes the empirical content of the finding.

The thesis contributes empirically by extending the horizon of the existing work on private capital and IPO governance through the post-pandemic cycle, the period in which the dual-class share of new listings rose from a minority practice to an increasingly prominent feature of American equity markets, and theoretically by reframing the relationship between private capital availability and IPO governance around the distinction between the stock and the flow of capital. The substantive content of these contributions is developed in Chapter 5 and summarized below alongside the limitations within which they hold.

The associative nature of the research design, the finite-cluster structure of the inference, and the absence of a direct measure of the bargaining mediator impose limits on the strength of the claims that can be drawn from this evidence, and those limits are not rhetorical; the distance between a documented association and a causal claim is real, and the directions identified in Chapter 5 for closing that distance, through quasi-experimental identification, through direct measurement of the mediator, and through comparative international variation in regulatory regimes, remain open for the research program this thesis initiates rather than concludes.

Within those limits, the contribution of the thesis is twofold. The empirical contribution is the demonstration that the choice between the stock and the flow of private capital, between dry powder and realized deal volume, is a measurable determinant of inference in the empirical study of IPO governance, with the flow measure delivering the tighter estimate when both are forced to compete. The theoretical contribution is the framing of realized late-stage deal-making, rather than the latent accumulation of uncommitted capital, as a candidate upstream mechanism whose operation is consistent with the observed translation of private capital market conditions into the governance architecture of newly listed firms.

The implication, advanced as a conditional reading rather than as a settled claim, is that the governance choices observed at the moment of listing are shaped in part by conditions that originate outside the domain of the instruments built to constrain them. The corrective instruments deployed by institutional investors and by the public-market regulatory apparatus were calibrated for a separation of ownership and control framed as the manager's problem; the present evidence indicates that the locus of the modern governance question has shifted upstream, into a domain those instruments were not designed to reach. That, more than any single coefficient the thesis reports, is what the evidence supports.

References

- Aggarwal, D., Eldar, O., Hochberg, Y. V., & Litov, L. P. (2022). The rise of dual-class stock IPOs. *Journal of Financial Economics*, 144(1), 122–153.
<https://doi.org/10.1016/j.jfineco.2021.12.012>
- Altman, S. (2025, May 5). *Evolving OpenAI's structure*. OpenAI. <https://openai.com/index/evolving-our-structure/>
- Aran, Y., & Pollman, E. (2024). Ousted. *Theoretical Inquiries in Law*, 25(1), 231–263.
<https://doi.org/10.1515/til-2024-0010>
- Bebchuk, L. A., & Kastiel, K. (2017). The untenable case for perpetual dual-class stock. *Virginia Law Review*, 103(4), 585–631. <https://virginialawreview.org/articles/untenable-case-perpetual-dual-class-stock/>
- Berle, A. A., & Means, G. C. (1932). *The modern corporation and private property*. Macmillan.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.
<https://doi.org/10.1162/rest.90.3.414>
- Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2), 317–372. <https://doi.org/10.3368/jhr.50.2.317>
- Cremers, M., Lauterbach, B., & Pajuste, A. (2018). *The life-cycle of dual class firms* (ECGI Finance Working Paper No. 550/2018). European Corporate Governance Institute.
<https://doi.org/10.2139/ssrn.3062895>
- Doidge, C., Karolyi, G. A., & Stulz, R. M. (2017). The U.S. listing gap. *Journal of Financial Economics*, 123(3), 464–487. <https://doi.org/10.1016/j.jfineco.2016.12.002>
- Ewens, M., & Farre-Mensa, J. (2020). The deregulation of the private equity markets and the decline in IPOs. *The Review of Financial Studies*, 33(12), 5463–5509.
<https://doi.org/10.1093/rfs/hhaa053>
- Ewens, M., Nanda, R., & Rhodes-Kropf, M. (2018). Cost of experimentation and the evolution of venture capital. *Journal of Financial Economics*, 128(3), 422–442.
<https://doi.org/10.1016/j.jfineco.2018.03.001>
- Field, L. C., & Lowry, M. (2022). Bucking the trend: Why do IPOs choose controversial governance structures and why do investors let them? *Journal of Financial Economics*, 146(1), 27–54.
<https://doi.org/10.1016/j.jfineco.2022.06.004>
- Friedman, A. M., McCormick, R., Rutherford, A. M., & Zivnuska, R. (2017, August 5). S&P and FTSE Russell on exclusion of companies with multi-class shares. *Harvard Law School Forum on Corporate Governance*. <https://corpgov.law.harvard.edu/2017/08/05/sp-and-ftse-russell-on-exclusion-of-companies-with-multi-class-shares/>
- Gao, X., Ritter, J. R., & Zhu, Z. (2013). Where have all the IPOs gone? *Journal of Financial and Quantitative Analysis*, 48(6), 1663–1692. <https://doi.org/10.1017/S0022109014000015>
- Gompers, P. A., Ishii, J., & Metrick, A. (2010). Extreme governance: An analysis of dual-class firms in the United States. *The Review of Financial Studies*, 23(3), 1051–1088.
<https://doi.org/10.1093/rfs/hhp024>

- Goshen, Z., & Hamdani, A. (2016). Corporate control and idiosyncratic vision. *The Yale Law Journal*, 125(3), 560–617. <https://yalelawjournal.org/article/corporate-control-and-idiosyncratic-vision>
- Jackson, R. J., Jr. (2018, February 15). *Perpetual dual-class stock: The case against corporate royalty* [Speech]. U.S. Securities and Exchange Commission. <https://www.sec.gov/news/speech/perpetual-dual-class-stock-case-against-corporate-royalty>
- Jordan, B. D., Kim, S., & Liu, M. H. (2016). Growth opportunities, short-term market pressure, and dual-class share structure. *Journal of Corporate Finance*, 41, 304–328. <https://doi.org/10.1016/j.jcorpfin.2016.10.003>
- Kim, H., & Michaely, R. (2019). *Sticking around too long? Dynamics of the benefits of dual-class voting* (ECGI Finance Working Paper No. 590/2019). European Corporate Governance Institute. <https://doi.org/10.2139/ssrn.3145209>
- Kwon, S., Lowry, M., & Qian, Y. (2020). Mutual fund investments in private firms. *Journal of Financial Economics*, 136(2), 407–443. <https://doi.org/10.1016/j.jfineco.2019.10.003>
- Loughran, T., & Ritter, J. R. (2004). Why has IPO underpricing changed over time? *Financial Management*, 33(3), 5–37. <https://site.warrington.ufl.edu/ritter/files/2015/06/Why-Has-IPO-Underpricing-Changed-Over-Time-2004.pdf>
- Mosco, G. D., & Felicetti, R. (2024). Gli assetti delle imprese che sviluppano l'intelligenza artificiale e il campanello dell'allarme OpenAI. *Giurisprudenza Commerciale*, 51(6), 1205–1216.
- PitchBook Data, Inc. (n.d.). *PitchBook Deals Screener* [Data platform]. PitchBook. <https://pitchbook.com/>
- PitchBook Data, Inc., & National Venture Capital Association. (2017). *4Q 2016 PitchBook-NVCA Venture Monitor*. PitchBook. <https://pitchbook.com/news/reports/4q-2016-pitchbook-nvca-venture-monitor>
- PitchBook Data, Inc., & National Venture Capital Association. (2024). *Q4 2023 PitchBook-NVCA Venture Monitor*. PitchBook. <https://pitchbook.com/news/reports/q4-2023-pitchbook-nvca-venture-monitor>
- PitchBook Data, Inc., & National Venture Capital Association. (2026). *Q4 2025 PitchBook-NVCA Venture Monitor*. PitchBook. <https://pitchbook.com/news/reports/q4-2025-pitchbook-nvca-venture-monitor>
- Ritter, J. R. (2026). *Initial public offerings: Updated statistics* [Dataset]. University of Florida. <https://site.warrington.ufl.edu/ritter/ipo-data/>
- Smart, S. B., & Zutter, C. J. (2003). Control as a motivation for underpricing: A comparison of dual and single-class IPOs. *Journal of Financial Economics*, 69(1), 85–110. [https://doi.org/10.1016/S0304-405X\(03\)00109-0](https://doi.org/10.1016/S0304-405X(03)00109-0)

Appendix

Figure A1

The Bargaining Power Channel From Late-Stage Private Capital to Dual-Class Adoption at IPO



Note. Solid arrows represent the hypothesized channel tested in this study (H1). Bargaining power is a theorized mediator, not directly measured.

Source. Author's elaboration based on Aggarwal et al. (2022); Ewens & Farre-Mensa (2020); Kwon et al. (2020).

Table A1*Sample Attrition Process*

Step	Observations
Ritter full IPO dataset (1975-2025)	16,030
Restricted to 2000-2023, non-ADR, non-closed-end, valid CRSP PERMNO	5,042
After exclusion scoring (SPACs, banks/S&Ls, REITs, LPs, etc.)	3,504
After dropping missing DualClass	3,480
After dropping missing SIC code	3,325
After dropping missing Age	3,250
After estimation exclusions (perfect separation in logit)	3,202
After dropping missing FirmSize (Compustat)	3,199

Note. Final analytical sample. The reduction from 3,250 to 3,202 reflects 48 observations lost to perfect separation of certain 2-digit SIC groups in the logistic estimation. The further reduction to 3,199 reflects 3 missing observations on total assets in Compustat, required for the FirmSize control.

Source. Author's elaboration based on Ritter (2026) and WRDS/Compustat.

Table A2
Variable Summary

Variable	Definition	Measurement	Scale	Source
DualClass	= 1 if the IPO adopts a dual-class share structure	S-1 filing classification	Binary	Ritter (2026)
LateStageVol	Aggregate annual volume of U.S. late-stage VC deals (Series C+)	Deal value in \$B	Continuous	PitchBook / NVCA VM
DryPowder_t1	Aggregate U.S. VC dry powder at 12/31 of year t-1	Capital overhang in \$B	Continuous	NVCA VM (Q4 2025)
LateStageCount	Number of U.S. late-stage VC deals per year	Deal count	Continuous	PitchBook / NVCA VM
VCBacked	= 1 if the firm received VC funding pre-IPO	VC-backed indicator	Binary	Ritter (2026)
TechInd	= 1 if the firm operates in a technology sector	36 SIC codes (Loughran & Ritter, 2004)	Binary	Loughran and Ritter (2004) / Ritter (2026)
Age	Years between founding date and IPO year	IPO year – Founding year	Continuous	Ritter (2026)
FirmSize	Natural logarithm of total assets	ln(Total Assets)	Continuous	WRDS/Compustat
ROA	Net income divided by total assets	Net Income / Total Assets	Continuous	WRDS/Compustat
Year	Calendar year of IPO	Linear time trend	Continuous	Ritter (2026)

Note. Industry fixed effects (2-digit SIC codes) are included in all specifications but are not listed as individual variables. LateStageCount is used in robustness checks only. DryPowder_t1 is included in baseline specifications for comparability but is not the variable on which the hypothesis test rests.

Table A3*Descriptive Statistics (N = 3,250)*

Variable	Mean	SD	Min	Max
DualClass	.177	.382	0	1
LateStageVol (\$B)	44.74	44.21	9.89	159.59
DryPowder_t1 (\$B)	98.44	50.90	50.75	281.27
VCBacked	.506	.500	0	1
TechInd	.208	.406	0	1
Age	17.70	23.81	0	179
FirmSize (ln AT)	5.53	1.71	−0.68	12.81
ROA (NI/AT)	−0.19	0.74	−28.74	0.86

Note. FirmSize and ROA have N = 3,247 due to 3 missing observations on Compustat. Descriptive statistics are computed on the pre-estimation sample.

Source. Author's elaboration based on Ritter (2026) and WRDS/Compustat.

Table A4*Logistic Regression Results - Progressive Specifications (Standard Errors Clustered by Year)*

Variable	M1	M2	M3	M4 (Main)
LateStageVol	0.0114*** (0.0018)	0.0135*** (0.0021)	0.0043*** (0.0009)	0.0038*** (0.0008)
VCBacked	-0.5815*** (0.1858)	-0.3595* (0.2099)	-0.3857* (0.1980)	-0.4934*** (0.1898)
TechInd	-0.1790 (0.1317)	-0.5485*** (0.2055)	-0.3327* (0.1756)	-0.3320** (0.1635)
Age	0.0030 (0.0025)	0.0026 (0.0030)	0.0028 (0.0032)	-0.0039 (0.0037)
FirmSize	—	—	—	0.2278*** (0.0669)
Year	—	—	0.0928*** (0.0109)	0.0940*** (0.0102)
Industry FE	No	Yes	Yes	Yes
Clustered SE (year)	Yes	Yes	Yes	Yes
N	3,250	3,202	3,202	3,199
Pseudo R ²	.063	.155	.178	.192
Log-likelihood	-1,422.50	-1,274.44	-1,240.34	-1,217.87

Note. Year-level clustered standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. The sample reduction from 3,250 to 3,202 in Model 2 is due to perfect separation of certain 2-digit SIC groups. The further reduction to 3,199 in Model 4 reflects 3 missing observations on FirmSize in Compustat.

Source. Author's elaboration. Stata 18 output.

Table A5*Average Marginal Effects - Model 4 (Standard Errors Clustered by Year)*

Variable	dy/dx	SE	z	p
LateStageVol	+0.000448	0.0000952	4.71	< .001
VCBacked	−0.05847	0.02253	−2.60	.009
TechInd	−0.03935	0.01912	−2.06	.040
Age	−0.00047	0.00043	−1.08	.278
FirmSize	+0.02699	0.00741	3.64	< .001
Year	+0.01114	0.00119	9.37	< .001

Note. N = 3,199. Average marginal effects computed via the delta method. Standard errors clustered at the year level.

Table A6*DryPowder_t1 versus LateStageVol - Full Specification (Standard Errors Clustered by Year)*

Variable	DP Only	LSV Only (= M4)	Both
DryPowder_t1	0.0037** (0.0018)	—	0.0027* (0.0015)
LateStageVol	—	0.0038*** (0.0008)	0.0027** (0.0011)
N	3,199	3,199	3,199
Pseudo R ²	.192	.192	.193

Note. All specifications include VCBacked, TechInd, Age, FirmSize, industry fixed effects (2-digit SIC), linear time trend, and standard errors clustered at the year level. Logistic coefficients reported. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A7a*Within-Design Robustness (Standard Errors Clustered by Year)*

Variable	M4 (Main)	R1 Count	R2 LPM	R3 +ROA	R4 Non-VC	R5 Non-Tech
<i>LateStageVol</i>	0.0038***	—	0.00095***	0.0039***	0.0027**	0.0030***
<i>LateStageCount</i>	—	0.00024***	—	—	—	—
<i>VCBacked</i>	−0.4934***	−0.4805**	−0.0569**	−0.4884***	—	−0.3347
<i>TechInd</i>	−0.3320**	−0.3378**	−0.0522**	−0.3376**	0.0673	—
<i>Age</i>	−0.0039	−0.0039	−0.0006	−0.0041	−0.0023	−0.0027
<i>FirmSize</i>	0.2278***	0.2334***	0.0309***	0.2523***	0.1541***	0.2193***
<i>ROA</i>	—	—	—	−0.1724	—	—
<i>Year</i>	0.0940***	0.0655***	0.0090***	0.0916***	0.0748***	0.1027***
<i>N</i>	3,199	3,199	3,247	3,199	1,541	2,522
<i>Pseudo R²</i>	.192	.192	.179	.194	.131	.195

Note. Standard errors clustered at the year level. All specifications include industry fixed effects (2-digit SIC) and a linear time trend. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. R2 reports linear probability model (OLS) coefficients; R-squared replaces Pseudo R^2 ($N = 3,247$ because OLS does not drop the perfect-separation observations). R4 excludes VCBacked as a regressor since the subsample consists exclusively of non-VC-backed firms; R5 excludes TechInd for the same reason. The structural-robustness block (R6a, R6b, R6c) is reported in Table A7b.

Source. Author's elaboration. Stata 18 output.

Table A7b*Structural Robustness - Alternative Time Trends (Standard Errors Clustered by Year)*

Variable	M4 (Main)	R6a Quadratic	R6b Piecewise	R6c Excl. 2020–21
<i>LateStageVol</i>	0.0038***	0.0038***	0.0017	0.0081***
<i>VCBacked</i>	−0.4934***	−0.4934***	−0.4902***	−0.6417***
<i>TechInd</i>	−0.3320**	−0.3320**	−0.3375**	−0.3590*
<i>Age</i>	−0.0039	−0.0039	−0.0040	−0.0026
<i>FirmSize</i>	0.2278***	0.2278***	0.2394***	0.2419***
<i>Year</i>	0.0940***	0.0940***	0.0526***	0.0846***
<i>Year_post2010</i>	—	—	0.0824***	—
<i>N</i>	3,199	3,199	3,199	2,683
<i>Pseudo R²</i>	.192	.192	.193	.174

Note. Standard errors clustered at the year level. All specifications include industry fixed effects (2-digit SIC). *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$. R6a's quadratic term (Year^2) was dropped by the maximum-likelihood routine due to empirical collinearity with Year over the 24-year window; reported coefficients therefore coincide with M4. R6b adds a piecewise-linear post-2010 trend (Year_post2010); under this specification LateStageVol is not statistically significant at conventional levels ($p = 0.110$) and the post-2010 trend term absorbs the cyclical variation that the linear trend leaves in the residual. R6c excludes the boom years 2020 and 2021 from the sample; the within-design checks (R1-R5) are reported in Table A7a.

Source. Author's elaboration. Stata 18 output.

Table A7*Inferential Diagnostics - Model 4*

Test	Statistic	Implication
Area under ROC curve	0.795	Acceptable discriminatory capacity
Wild cluster bootstrap, LateStageVol	$p = 0.014$, $z = 2.09$	Significant at 5% under finite-sample correction

Note. Both diagnostics applied to the Model 4 specification reported in Table A4. The wild cluster bootstrap follows Cameron, Gelbach, and Miller (2008) with 9,999 Rademacher replications and year-level bootstrap clustering, using a fixed random seed for reproducibility. ROC and bootstrap address inferential properties (model discrimination and small-cluster inference, respectively); they do not address the structural identification limitations discussed in Sections 3.4.2, 3.5.1, and 3.7.5.

Source. Author's elaboration. Stata 18 output.

Appendix materials: supplementary replication files, Stata do-files, and additional robustness checks are available upon request. The PitchBook Deals Screener filtering criteria and the discrepancy check sheet are filed as supplementary replication materials.

Acknowledgements

This thesis, which in its pages fixes the point of arrival of a four-year undergraduate path, carries with it something that cannot be reduced to paragraphs and citations: the fears that accompanied it, the anxieties that punctuated its passages, the subtle doubt that at moments called into question its very possibility. If these pages exist today, it is because those fears and those anxieties were not crossed alone; and the fact of having crossed them, more than the result that these pages seek to represent, constitutes the thing that truly matters.

To my parents, who allowed me to be the first in our family to walk this road, I owe more than can be said. The freedom to choose an undergraduate path distributed across Rome, Beijing, and Washington, to build an education that no one before me in our family history had been able to imagine, rests entirely on their silent sacrifice; a sacrifice that they never sought to make visible, and one that I began to understand only to the extent that, as I grew older, I learned to recognize its real weight. If there exists a continuity between the boy who left Rome at eighteen and the one who today submits this thesis, that continuity is theirs, theirs alone.

To Ludovica, whose constant presence in every phase of this journey represented the steady point that allowed me to carry it to completion, I owe the certainty of never having had to face alone the moments in which determination alone would not have sufficed. Her capacity to keep me anchored when deadlines accumulated, to listen when words had not yet found a form, to be there simply when being there was all that was needed, is inscribed in every page that precedes this one, even if it appears in none.

These pages close nothing; they open only a subsequent phase in which the fears that accompanied these three years will transform into others, and the anxieties into others still. But the certainty of having already crossed them once, together, is what these pages carry with them as their most enduring contribution.