



UNIVERSITÀ
DI PAVIA

**Dipartimento di Scienze Economiche e
Aziendali**

**MASTER IN ECONOMICS,
DEVELOPMENT AND
INNOVATION**

**The effect of Venture Capital Support on exit
trajectories: evidence from Cybersecurity Startups
(2009-2024)**

Relatore:

Chiar.mo Prof. Roberto Fontana

**Tesi di Laurea
di Chiara Scrascia**

Matr. n. 551492

Anno Accademico 2025-2026

*Ai nuovi obiettivi,
Alle nuove opportunità*

ENGLISH ABSTRACT

Today cybersecurity is a strategic and dynamic sector, sitting at the intersection of commercial growth, technological competition and geopolitical relevance. The entrepreneurship literature offers an established model for explaining startup survival, the role of founders and the function of venture capital, but this model was designed for tech startups in general, not specifically for cybersecurity, which has distinctive characteristics of its own: competitive fragmentation, exits dominated by acquisitions rather than IPOs, a "cockroach" model based on capital efficiency, and periodic regulatory shocks.

This thesis tests empirically whether, and how, these mechanisms apply to the cybersecurity sector, tracing a startup's full trajectory: from access to financing, through survival, to the way it eventually leaves the market.

The study draws on an original dataset of 1,227 European cybersecurity startups (Crunchbase, 2009–2024) and adopts a two-stage empirical strategy. In the first stage, a logit model estimates the probability of securing equity financing as a function of founding-team composition, founder background, and sector positioning. In the second stage, a Cox survival model estimates the determinants of survival, using the financing propensity predicted in the first stage as an instrument to address the endogeneity between financing and firm performance. The analysis concludes with a competing-risks extension, which separates market exit into two mutually exclusive, cause-specific outcomes: operational failure and acquisition.

The results paint a more nuanced picture than the general entrepreneurship literature suggests. Founding-team composition is the strongest predictor of access to capital, while individual founders' credentials show no significant effect. Survival, conditional on financing, is not explained by firm age: the hazard of exit rises in an almost linear fashion over time, without the early peak predicted by the classical "liability of newness", and is instead explained by structural factors such as firm size and sector cluster.

The thesis's most original contribution comes from disaggregating failure and acquisition, and it points in an unexpected direction relative to the literature that frames venture capital as an accelerant of polarised outcomes: financing propensity proves not to be a significant determinant of either the timing or the channel of market exit. Firm size, rather, drives the mode of exit. Larger firms do not exit the market more often, but they exit in a systematically different way, almost exclusively through acquisition. Operational failure, by contrast, is a rare phenomenon overall and is concentrated mainly among medium-sized firms (201 to 1,000 employees), while both smaller and larger firms show a near-zero failure risk. Sector adds a further layer of insight: Core cybersecurity firms show a significantly higher acquisition risk, not a higher failure risk, a sign of a particularly active acquisition market in this technology segment.

These findings carry practical implications for both investors and policymakers: risk assessment in the cybersecurity sector should rest on structural indicators, such as firm size and sector, rather than on firm age or financing propensity alone; and the rising level of M&A activity in the core cybersecurity segment raises competition-policy questions that merit regulatory attention.

ITALIAN ABSTRACT

La cybersecurity è oggi un settore strategico e dinamico, a metà tra crescita commerciale, competizione tecnologica e rilevanza geopolitica. La letteratura sull'imprenditorialità offre un modello consolidato per spiegare sopravvivenza delle startup, ruolo dei fondatori e funzione del venture capital — ma pensato per le tech startup in generale, non per la cybersecurity, che ha caratteristiche proprie: frammentazione competitiva, uscite dominate da acquisizioni più che IPO, modello "cockroach" basato sull'efficienza del capitale, e shock regolatori periodici.

Questa tesi verifica empiricamente se e come tali meccanismi si applicano al settore cybersecurity, seguendo l'intero percorso di una startup: dall'accesso al finanziamento, alla sopravvivenza, fino alla modalità di uscita dal mercato.

Lo studio si basa su un dataset originale di 1.227 startup cybersecurity europee (Crunchbase, 2009–2024) e adotta una strategia empirica a due stadi. Nel primo stadio, un modello logit stima la probabilità di ottenere un finanziamento equity in funzione della composizione del team fondatore, del background dei fondatori e del posizionamento settoriale. Nel secondo stadio, un modello di sopravvivenza di Cox stima i determinanti della sopravvivenza, usando la propensione al finanziamento prevista al primo stadio come strumento per limitare il problema di endogeneità tra finanziamento e performance. L'analisi si conclude con un'estensione a rischi competitivi, che distingue due modalità di uscita dal mercato, tra loro alternative: il fallimento operativo e l'acquisizione.

I risultati restituiscono un quadro più articolato di quanto suggerisca la letteratura generale sull'imprenditorialità. La composizione del team fondatore è il predittore più solido dell'accesso al capitale, mentre le credenziali individuali dei singoli fondatori non hanno un effetto significativo. La sopravvivenza, una volta ottenuto il finanziamento, non dipende dall'età dell'impresa: il rischio di uscita cresce in modo quasi lineare nel tempo, senza il picco iniziale previsto dalla classica "liability of newness", ed è spiegato piuttosto da fattori strutturali come la dimensione aziendale e il cluster settoriale.

Il contributo più originale della tesi riguarda proprio la scomposizione tra fallimento e acquisizione, e va in una direzione inattesa rispetto alla letteratura che descrive il capitale di rischio come acceleratore di esiti polarizzati: la propensione al finanziamento non risulta un determinante significativo né dei tempi né del canale di uscita dal mercato. A guidare la modalità di uscita è piuttosto la dimensione aziendale: le imprese più grandi non escono dal mercato più spesso, ma escono in modo sistematicamente diverso, quasi sempre per acquisizione. Il fallimento operativo, al contrario, è un fenomeno nel complesso raro e concentrato soprattutto tra le imprese di dimensione media (tra 201 e 1.000 dipendenti), mentre le imprese più piccole e quelle più grandi mostrano un rischio di fallimento quasi nullo. Il settore aggiunge un ulteriore livello di lettura: le imprese Core cybersecurity mostrano un rischio di acquisizione significativamente più elevato, non un rischio di fallimento più alto, segno di un mercato di acquisizioni particolarmente attivo in questo segmento tecnologico.

Questi risultati hanno implicazioni concrete sia per gli investitori sia per i policy maker: la valutazione del rischio nel settore cybersecurity dovrebbe basarsi su indicatori strutturali, come la dimensione dell'impresa e il settore di appartenenza, più che sulla sola età anagrafica o sulla propensione al finanziamento; inoltre, la crescente attività di M&A nel segmento core cybersecurity solleva questioni di politica della concorrenza che meritano attenzione regolatoria.

ACKNOWLEDGEMENTS

Desidero anzitutto rivolgere un sentito ringraziamento al mio relatore, Prof. Roberto Fontana, per il prezioso supporto, la disponibilità e la competenza con cui ha seguito lo sviluppo di questo lavoro. Nonostante le nostre interazioni si siano svolte a distanza, non è mai venuta meno una presenza costante e attenta, fatta di consigli puntuali e osservazioni critiche. La sua guida è stata fondamentale non solo per la stesura di questa tesi, ma anche per la crescita del mio metodo di lavoro e del mio pensiero critico. Lo ringrazio per la pazienza dimostrata, per il tempo dedicatomi e per aver sempre creduto nella riuscita di questo percorso.

Raggiungere questo nuovo traguardo segna la “fine” di un percorso scolastico e l’inizio di un periodo più incerto, che porterà più ansie, preoccupazioni ma che alla fine darà sempre grandi soddisfazioni.

Questi ringraziamenti sono un pò diversi dal solito perchè prima di voler ringraziare tutte le persone a me care, voglio auto-ringraziarmi. Sono riuscita a raggiungere tutti gli obiettivi che mi sono sempre prefissata, sono fiera di me stessa, dei miei sforzi. Sono grata di essere qui oggi, insieme a tutte le persone (o parte) che hanno fatto parte della mia crescita. Chi mi conosce bene sa che questi mesi non sono stati semplici, motivo in più per cui mi sento così soddisfatta oggi, al di là del voto.

Esattamente 2 anni fa ho iniziato un percorso che non sapevo se sarei riuscita a portare a termine: tra la mole di studio e la lingua inglese. All’inizio non è stato facile, ho dovuto impegnarmi, migliorandomi giorno per giorno per essere alla pari dei miei compagni e allo stesso tempo studiare meticolosamente tutto. Ma con più pazienza e sacrificio ce l’ho fatta. Ad oggi posso dire che ho il livello “intermedio” di Spagnolo/colombiano che posso aggiungere alle mie competenze linguistiche.

Come sono riuscita ad affrontare tutto è sempre grazie alla mia famiglia che mi è sempre stata accanto, incoraggiandomi e supportandomi, anche nei miei ultimi 9 mesi lontano da loro. Non mi sono mai sentita sola e sapevo che per qualsiasi cosa, per condividere traguardi, lamentarmi, o chiedere aiuto loro erano sempre lì. Sono sempre stati, e sempre saranno, la mia forza e quel posto sicuro. Sono così felice di avere una famiglia tanto unita con me, che mi insegna i veri valori della vita e mi appoggia nelle mie decisioni, e mostrandomi sempre strade da poter intraprendere.

Ai miei genitori e mia sorella, il mio ringraziamento più puro: vi voglio bene e grazie per non avermi fatto mai mancare nessuna opportunità di crescita.

Alla mia Alice, sempre lì, da più di 10 anni ormai: abbiamo vissuto l'una dell'altra i nostri traguardi, sapendo tutto ciò che significavano. Conosci i miei punti di debolezza e mi hai insegnato a trasformarli in punti di forza; sei sempre stata affianco a me, non lasciandomi mai sola. Sono davvero fortunata ad averti nella mia vita.

I miei momenti di follia vera, di pianti isterici e di ansia di quando ero a Tarragona erano sempre calmati dalla dolcezza di Jordan, che è entrato così lentamente nella mia routine, standomi sempre al mio fianco. Aquí no puedo hacerte la dedicatoria de amor, pero quiero tener presente decirte y recordarte lo especial que eres para mí. No ha sido fácil afrontar todo este período de estrés, pero contigo a mi lado todo parecía más sencillo. He confiado siempre en el amor porque veo a mis padres todos los días, que entre peleas y períodos de serenidad, siempre están el uno al lado del otro; y así es como me imagino una vida contigo, quiero vivirte todos los días y saber que pase lo que pase, tanto a mí como a ti, el uno siempre podrá contar con el otro, tal como me lo has demostrado en estos meses juntos.

La mia tranquillità a Tarragona è stata segnata anche da due personcine trovate per caso: ad Elena e Jasmin, che mi avete riempito le giornate con risate e periodi di serenità, sono felice di avervi conosciute perchè con voi mi sono sentita a casa stando lontana da casa.

Alle mie amiche, ormai da più di 5 anni, a Giulia, Chiara, Bea, Giorgia e Martina sempre lì pronte per stare insieme anche solo per un caffè al volo. Vi auguro sempre il meglio, non solo accademicamente, siete persone di cuore e nonostante la distanza, abbiamo sempre trovato un modo per sapere come sta l'altra.

Un bene infinito a Greta, Elisa, Rebecca che tra chiamate perse, 1000 messaggi, podcast di audio non mi hanno mai fatto sentire sola.

Ai miei compagni di corso e in particolare a Roberta e Livia. Tra ore di studio disperato e tra risate siamo riuscite in questa grande impresa che all'inizio sembrava infinita.

Ovviamente ho già il mio obiettivo: lavorare in Spagna, apprendere meglio la lingua, poter sfruttare tutto ciò che in questi anni ho imparato e mi hanno insegnato. Un ultimissimo ringraziamento va a Cristina Dumitriu: con te ho imparato molto, da come individuare un problema a come affrontarlo efficientemente, ho visto in te una figura che mi piacerebbe raggiungere; grazie per avermi insegnato tanto in quei pochi mesi in Mamenteive.

Grazie a tutti coloro che oggi sono qui con me per festeggiare questo giorno.

Desidero sia speciale, pieno di buoni ricordi.

TABLE OF CONTENTS

LIST OF TABLES -----	14
LIST OF FIGURES -----	15
CHAPTER 1: INTRODUCTION -----	17
1.1 MOTIVATION AND STRUCTURAL FEATURES -----	17
1.2 PROBLEM STATEMENT AND RESEARCH QUESTIONS -----	19
1.3 OBJECTIVES AND EXPECTED CONTRIBUTION -----	20
1.4 STRUCTURE OF THE THESIS -----	21
CHAPTER 2: REVIEW OF LITERATURE -----	25
2.1 STARTUP LIFECYCLE AND THE “VALLEY OF DEATH” -----	25
2.1.1 <i>The concept of survival</i> -----	25
2.1.2 <i>Startup lifecycle</i> -----	26
2.1.3 <i>The Valley of Death</i> -----	29
2.2 THE RELATIONSHIP BETWEEN FOUNDER BACKGROUND AND VC CAPITAL INVESTMENT -----	31
2.2.1. <i>The high-opportunity-cost founders as a double- edged sword of capital</i> -----	31
2.2.2 <i>Macroeconomic aspects: hot and cold markets</i> -----	32
2.3 THE RELATIONSHIP BETWEEN THE VC CAPITAL INVESTMENT AND THE MODE OF EXIT -----	34
2.3.1 <i>“Cash-Out” and “Flame-Out”</i> -----	34
2.3.2 <i>The role of macroeconomic cycles</i> -----	35
2.4 SUMMARY AND THEORETICAL FRAMEWORK -----	36
CHAPTER 3: THE CONTEXT OF ANALYSIS: THE CYBERSECURITY INDUSTRY -----	39
3.1 INTRODUCTION -----	39
3.2 MARKET SIZE AND GROWTH TRENDS -----	39
3.3 STRUCTURAL CHARACTERISTICS OF THE CYBERSECURITY STARTUP ECOSYSTEM -----	40
3.3.1 <i>Market fragmentation</i> -----	40

3.3.2 <i>The "cockroach" versus the "unicorn" survival model</i>	43
3.3.3 <i>Venture Capital role and mode of exit</i>	44
3.3.4 <i>The nature of innovation</i>	46
3.4 THE VALLEY OF DEATH IN CYBERSECURITY VENTURES	47
3.5 OVERVIEW AND FINAL DISCUSSION	48
CHAPTER 4: DATA AND METHOD	51
4.1 INTRODUCTION AND GENERAL DESCRIPTION OF THE DATASET	51
4.2 CLUSTER VARIABLES	55
4.3 TEMPORAL DIMENSIONS OF THE DATASET	59
4.3.1 FOUNDATION PERIOD	59
4.3.2 <i>Closure trend</i>	61
4.3.3 <i>Entry cohorts</i>	62
4.3.4 <i>Lifespan and survival</i>	63
4.4 FIRMS AND FOUNDERS CHARACTERISTICS	65
4.4.1 <i>Solo Founders versus Founding Teams</i>	67
4.4.2 <i>Founders' educational background</i>	68
4.4.3 <i>Professional background</i>	71
4.4.4 <i>Founding rounds</i>	72
4.4.5 <i>Mode of exit</i>	74
4.5 OTHER VARIABLES	76
4.5.1 <i>Revenues</i>	76
4.5.2 <i>Geographical distribution</i>	77
4.5.3 <i>Company' size</i>	78
4.6 SUMMARIZE AND METHOD APPLIED	79
4.6.1 <i>Overview and econometric strategy</i>	79
4.6.2 <i>Method applied</i>	80
CHAPTER 5: EMPIRICAL ANALYSIS	83
5.1 OVERVIEW	83
5.2 LOGIT ESTIMATION	84
5.3 KAPLAN-MEIER ANALYSIS	88
5.4 SURVIVAL MODEL: COX PROPORTIONAL HAZARDS MODEL	93

5.4.1 Centred regression -----	95
5.4.2 Robustness check: uncentred regression-----	98
5.5 BASELINE FUNCTIONS -----	101
5.6 MODE OF EXIT-----	103
5.7 OVERVIEW OF FINDINGS AND DISCUSSION -----	110
5.7.1 Research question 1: determinants of equity financing-----	111
5.7.2 Research question 2: determinants of survival conditional on financing -----	113
5.7.3 Research question 3: firms' exit route (conditional on being funded) -----	115
5.8 FINAL DISCUSSION: CONFRONTING THE EVIDENCE WITH THE LITERATURE -	118
5.8.1 Organisational survival and the Valley of Death -----	118
5.8.2 Founder background as a signal to investors -----	120
5.8.3 VC capital as an accelerator of polarised outcomes -----	120
CHAPTER 6: CONCLUSION -----	123
6.1 SUMMARY -----	123
6.2 IMPLICATIONS -----	124
6.3 LIMITATIONS -----	127
BIBLIOGRAPHY -----	131

LIST OF TABLES

TABLE 3.1: MARKET STRUCTURE-----	41
TABLE 4.1: OVERVIEW OF DATASET VARIABLES -----	54
TABLE 4.2:MARKET FRAGMENTATION-----	59
TABLE 4.3: SUMMARY STATISTICS ON LIFESPAN AND SURVIVAL -----	63
TABLE 4.4: T-TEST: FIRM LIFESPAN BY STATUS -----	64
TABLE 4.5: MODE OF EXIT AMONG CLOSED FIRMS -----	74
TABLE 4.6: FOUNDING TYPE AMONG CLOSED FIRMS -----	75
TABLE 4.7: MODE OF EXIT BY FUNDING TYPE -----	76
TABLE 4.8: DISTRIBUTION OF REVENUE CATEGORIES-----	77
TABLE 4.9: DISTRIBUTION BY GEOGRAPHIC LOCATION -----	78
TABLE 4.10: DISTRIBUTION OF SIZE FIRMS' CATEGORIES -----	78
TABLE 5.1:LOGIT MODEL -----	86
TABLE 5.2: SURVIVAL-TIME DATA ANALYSIS-----	89
TABLE 5.3: LOG-RANK TEST RESULTS -----	89
TABLE 5.4: SUMMARY STATISTICS PREDICTED (EQ_DUMMY_HAT) -----	93
TABLE 5.5: CENTRED REGRESSION -----	96
TABLE 5.6: UNCENTRED REGRESSION -----	99
TABLE 5.7: CAUSE-SPECIFIC COX MODEL: MODE OF EXIT(ACQUISITION- SHUTDOWN) -----	108

LIST OF FIGURES

FIGURE 4.1: OVERVIEW OF COMPANIES' TYPE-----	56
FIGURE 4.2: DISTRIBUTION OF COMPANIES BY FOUNDATION YEAR -----	60
<i>FIGURE 4.3: CLOSURE TREND OF COMPANIES-----</i>	<i>61</i>
<i>FIGURE 4.4: STSET DURATION, WITHIN 6 YEARS THRESHOLD-----</i>	<i>65</i>
FIGURE 4.5: DISTRIBUTION OF SOLO FOUNDER VS FOUNDING TEAMS -----	67
FIGURE 4.6: EDUCATIONAL BACKGROUND OF TEAM FOUNDER VS SOLO FOUNDER -	70
FIGURE 4.7: PROFESSIONAL BACKGROUND OF TEAM VS SOLO FOUNDER-----	72
FIGURE 4.8: FOUNDING ROUNDS ACHIEVED BY COMPANIES -----	73
FIGURE 5.1: KM SURVIVAL ESTIMATES BY FUNDING STATUS (EQ_DUMMY) -----	90
FIGURE 5.2: KM SURVIVAL ESTIMATES BY TEAM COPOSITION (TEAM_DUMMY)----	91
FIGURE 5.3: KM ESTIMATES BY PROFESSIONAL BACKGROUND (PROF_BG)-----	92
FIGURE 5.4: KM ESTIMATES BY EDUCATIONAL BACKGROUND (EDU_BG) -----	92
<i>FIGURE 5.5: BASELINE SURVIVAL FUNCTION, ESTIMATED FOR THE AVERAGE FIRM--</i>	<i>101</i>
FIGURE 5.6: BASELINE CUMULATIVE HAZARD FUNCTION, ESTIMATED FOR THE AVERAGE FIRM -----	102
FIGURE 5.7: DISTRIBUTION OF MODE OF EXIT -----	105
FIGURE 5.8: MODE OF EXIT BY SIZE CATEGORY-----	106
FIGURE 5.9: MODE OF EXIT BY REVENUE CATEGORY -----	107
FIGURE 5.10: KM CONDITIONAL SURVIVAL BY FIRM SIZE -----	114
FIGURE 5.11: BANKRUPTCY RISK-----	117

CHAPTER 1: INTRODUCTION

1.1 Motivation and structural features

Cybersecurity has become one of the most consequential segments of the technology economy. It is also one of the most closely watched. Industry estimates place the global market at approximately USD 217.9 billion in 2021. That figure is projected to reach USD 345.4 billion by 2026, a compound annual growth rate of 9.7% (Moiseienko and Kiva, 2022).

This growth tracks the expansion of the Internet of Things. Connected devices worldwide numbered approximately 10.94 billion in 2020. That number is projected to reach 25.44 billion by 2030 (Moiseienko and Kiva, 2022). Each new device widens the attack surface. Cybersecurity firms exist to defend it.

The sector's importance is not merely commercial. Cybersecurity sits at the intersection of corporate strategy, national security, and geopolitical competition. Data breaches, ransomware attacks, and critical-infrastructure incidents carry costs that reach well beyond the firms directly affected. Governments increasingly treat cybersecurity capability as a matter of strategic autonomy, not simply a line in an IT budget. This has direct implications for how cybersecurity ventures are financed, and for who is willing to acquire them.

The sector also sustains an unusual intensity of competition and technological change, simultaneously. Threats evolve continuously. Defensive technology must keep pace, or move faster still, simply to remain relevant. Competitive advantage is therefore rarely permanent. A leading solution today can be overtaken within a short window, by a new class of attack, or by a competitor's next release. Innovation is not a discretionary source of growth in this sector. It is a condition of survival. This dynamic sets the motivating context for the thesis. It is not tested directly in the empirical analysis, which relies on observable firm- and founder-level

characteristics rather than on measures of technological pace or innovation intensity.

This dynamic raises the question that motivates this thesis. A small number of large, established players, built on decades of accumulated infrastructure, customer relationships, and platform integration, dominate large segments of the cybersecurity market. Yet the industry's technological frontier is pushed forward disproportionately by small and medium-sized startups, operating with far fewer resources. This asymmetry is worth understanding in its own right.

Three structural features stand out.

1. First, the market is highly fragmented: no single security platform is universally applicable, and startups instead cluster around narrow, specialised "attack vectors" (Section 3.3.1).
2. Second, cybersecurity ventures tend to follow a "cockroach" survival logic, prioritising capital efficiency and resilience over the rapid, loss-financed scaling typical of consumer-tech "unicorns" (Section 3.3.2).
3. Third, exit runs almost exclusively through acquisition rather than public listing, driven by two distinct buyer types: platform providers absorbing niche capabilities directly into their core products, and telecoms carriers bundling security into managed infrastructure (Section 3.3.3).

These features shape how cybersecurity startups access financing, how long they survive, and how they eventually leave the market.

Startups in this sector rarely aim to displace incumbents outright. Prolonged, direct competition is not their usual goal. Their role is narrower, and more specific: i) first, they develop a defensive capability quickly, in response to an emerging threat; ii) second, they bring that capability to market before an incumbent can build it internally. For many startups, the intended endpoint is not independent survival. It is absorption into a larger platform, through acquisition, at the point where their technology is most valuable to a consolidator. This dynamic between incumbents

and startups sets the broader context for this thesis. It is not, however, the object of empirical analysis here.

Instead, this thesis turns to a narrower set of questions: how these startups secure financing in the first place, what shapes their chances of survival once funded, and how they eventually leave the market. These questions are set out formally in Section 1.2.

1.2 Problem statement and research questions

The academic literature, reviewed in Chapter 2, offers a well-developed account of the relationship between organisational survival, founder background, and venture capital. This account converges on a single, sequential narrative: survival depends on a venture's stage of development; founder background functions as a strategic signal shaping access to financing; and venture capital acts as an accelerant, polarising a venture's eventual outcome towards either a successful cash-out or an outright flame-out.

This framework is deliberately sector-agnostic. It draws on general entrepreneurial and corporate finance literature, developed for high-technology ventures at large rather than for any one sector. Several of its core mechanisms, however, plausibly manifest with different intensity within cybersecurity specifically. The sector's competitive structure is fragmented and narrowly segmented. Its exit landscape is bimodal, dominated by acquisition rather than public listing. It sustains a distinctive "cockroach" survival culture, which rewards capital efficiency over rapid scaling. And it remains exposed to periodic regulatory shocks, which reshape acquisition activity independently of firm performance. The general literature does not directly account for any of these features.

This creates a specific empirical gap. A general theory of founder background and venture capital, applied without adaptation, risks missing the mechanisms that make cybersecurity ventures distinctive.

The dataset used in this thesis reflects this distinctiveness directly. It comprises 1,227 privately held cybersecurity startups, based in Europe, with no publicly listed firms and no IPO exits recorded. The "cash-out" pathway central to the general theoretical framework can therefore only be tested through acquisition, not through public listing. This constraint is itself informative about how liquidity events actually occur in this sector.

This thesis is structured around three research questions, addressed sequentially:

RQ1. What variables influence the likelihood of receiving funding for cybersecurity startups?

RQ2. Conditional on being funded, what are the determinants of survival of cybersecurity startups?

RQ3. Conditional on being funded, in what way do firms exit the market, whether through acquisition or failure?

Each question addresses a distinct stage of a venture's trajectory, and each is investigated through a different empirical tool.

RQ1 addresses the financing threshold a venture must cross. It is investigated through a logit model, estimating the probability of securing equity financing as a function of founder- and firm-level characteristics.

RQ2 addresses what happens to survival once that threshold has been crossed. It is investigated first descriptively, through Kaplan-Meier survival estimates, and then through a Cox proportional hazards model, which isolates the effect of financing propensity on exit risk, net of other structural factors.

RQ3 addresses how, and not merely whether, the venture eventually leaves the market. It is investigated through a competing-risks extension of the Cox model, separating acquisition from shutdown as distinct, mutually exclusive outcomes.

1.3 Objectives and expected contribution

This thesis aims to make three contributions.

First, an empirical contribution. It applies a two-stage empirical strategy. A first-stage logit model estimates financing probability. A second-stage Cox proportional hazards model estimates survival. This second stage is further extended through a competing-risks framework, distinguishing acquisition from shutdown as separate outcomes. The strategy is applied to a sample of 1,227 cybersecurity startups. The predicted financing propensity from the first stage is used as an instrument in the second. This addresses the endogeneity that arises from treating financing status as directly exogenous to survival. Few studies apply this two-stage design specifically to the cybersecurity sector.

Second, a theoretical contribution. This thesis tests a general framework of founder background, financing, and exit against sector-specific evidence. It assesses which mechanisms generalise to cybersecurity, and which require revision. Particular attention is paid to the role of firm age. The general liability-of-newness literature makes a specific, testable prediction about the shape of exit risk over time. Whether this prediction holds within a sector marked by a distinctive "cockroach" survival culture, rather than by rapid, loss-financed scaling, is an open empirical question. The thesis is positioned to speak directly to this question.

Third, a practical contribution. The findings are intended to be of direct relevance to three audiences. Investors may draw on them to assess the risk profile of cybersecurity ventures. Founders may draw on them to navigate the financing process more effectively. Policy makers may draw on them to design support schemes, or to monitor competitive dynamics in a sector of growing strategic importance. Each audience faces a different decision, but each decision depends on the same underlying question: what actually drives financing, survival, and exit in this sector.

1.4 Structure of the thesis

The remainder of this thesis is organised as follows.

Chapter 2 reviews the general literature underpinning the empirical strategy. It opens with the concept of organisational survival as a time-dependent, rather than

static, condition, and examines the startup lifecycle and the Valley of Death, the phase of heightened hazard that follows founding. It then turns to founder background, discussing how professional and educational credentials, and founding-team composition, function as signals to investors under conditions of information asymmetry, moderated by the macroeconomic financing cycle. The chapter closes by examining the role of venture capital, not merely as a facilitator of survival but as an accelerant that polarises a venture's eventual trajectory towards either a successful cash-out or an outright flame-out. A concluding section draws these three strands together into a single theoretical framework and maps that framework onto the two-stage empirical strategy developed in later chapters.

Chapter 3 narrows this general framework to the cybersecurity sector specifically. It examines the market's size and growth trajectory, its fragmented and narrowly-segmented competitive structure, and the distinctive "cockroach" survival culture that rewards capital efficiency over rapid scaling, in contrast with the "unicorn" model more commonly associated with venture-backed technology firms. It also considers the nature of technological discontinuity within the sector, and how this shapes both competitive dynamics and acquisition activity. The chapter concludes by formally stating the three research questions that guide the empirical analysis.

Chapter 4 presents the dataset underlying the empirical analysis. It describes the sample of 1,227 cybersecurity startups, drawn from Crunchbase, and sets out the construction of the key variables used throughout the thesis. This includes the sector clustering scheme used to group firms into comparable segments, the construction of the survival-time variable, and the coding of exit outcomes into active, shutdown, and acquisition. Descriptive statistics are presented alongside a discussion of the dataset's coverage and limitations.

Chapter 5 presents the empirical analysis in full. It begins with a first-stage logit model, estimating the probability of securing equity financing as a function of founder- and firm-level characteristics. It then presents descriptive Kaplan-Meier survival estimates, before turning to a second-stage Cox proportional hazards model, in which the predicted financing propensity from the first stage is introduced

as an instrumented covariate. The chapter also derives and interprets the estimated baseline hazard function and closes with a competing-risks extension that disaggregates market exit into two distinct, cause-specific outcomes: acquisition and shutdown.

Chapter 6 discusses these results in relation to each research question in turn, before confronting the evidence with the theoretical framework developed in Chapter 2, strand by strand. It sets out the theoretical, managerial, and policy implications of the findings, and closes with a discussion of the limitations of the analysis and its data.

CHAPTER 2: REVIEW OF LITERATURE

2.1 Startup lifecycle and the “valley of death”

2.1.1 The concept of survival

Startups play a pivotal role in driving economic growth and innovation, rendering their survival crucial for a thriving economy (OECD, 2023). However, startup survival extends beyond macroeconomic impact.

Within the entrepreneurship literature, the concept of organisational survival is theoretically rooted in resource-based entrepreneurship theories (Esteve-Pérez, 2008), psychological entrepreneurship theories, and opportunity-based entrepreneurship theories (Simpeh, 2011; Hejazi & Seyyedamiri, 2024).

Conversely, the conceptualisation of firm failure is shaped by both cognitive limitations (Franco and Haase, 2010) and the specific “baggage” of professional pre-entry experience (Paik, 2014). Franco and Haase (2010) argue that structural vulnerability often stems from a self-serving bias. Founders' inability to objectively analyse failure results in stunted organisational learning and network constraints. Furthermore, Paik (2014) posits that prior venture capital experience can induce a counterproductive "cost of learning," which paradoxically leads to lower survival rates for subsequent ventures.

Hejazi and Nader Seyyedamiri (2024) categorise the study of survival into two primary approaches. Some scholars examine survival through the lens of ongoing organizational processes, while others evaluate it as a definitive, static outcome metric. At its core, survival denotes the continued existence of an enterprise over a specific timeline following its inception. It reflects a firm's capability to continue operating as a self-sustaining economic entity and its capacity to respond successfully to rapidly changing environmental conditions. Consequently, while rapid expansion is not a strictly necessary condition for long-term viability, survival itself represents the fundamental short-term objective to prevent failure, particularly for newly established ventures exposed to internal and external environmental threats.

2.1.2 Startup lifecycle

Analysing survival requires a developmental perspective, acknowledging that the factors influencing a firm's viability shift as it evolves. To model this evolution, the literature often evaluates firms based on their "stage of existence" (Agarwal & Gort, 2002), heavily weighting their "initial endowments." In econometric terms, this nascent stage is directly associated with an elevated baseline hazard rate. This mortality risk is not merely a static penalty for "newness", but rather a dynamic function of two opposing forces: the rapid obsolescence of founding endowments and the firm's capacity for "learning-by-doing".

A widely accepted conceptualisation of the firm lifecycle is Awasthi (2011) *Four-Stage Lifecycle Model*, which divides an enterprise's development into four sequential phases:

0. Pre-startup: the conceptualisation and planning phase.
1. Startup: the entry phase, characterised by active market learning and capability acquisition.
2. Survival: the critical transition phase. Many ventures stagnate or fail at this juncture; however, entrepreneurs who successfully navigate this stage reach a point where business activities stabilise.
3. Growth: the expansion phase, which enhances the probability of long-term viability.

The survival phase represents a critical juncture. Whilst many ventures tend to stagnate at this stage, entrepreneurs who navigate it seamlessly reach a phase of stabilisation and consolidation across all commercial operations.

Bridging the concepts of the "liability of newness" (Lyles et al., 2004) and the hazard rate (Audretsch & Mahmood, 1995), the literature highlights why this initial stage is so precarious. A startup earliest phase is characterised by a high probability of failure due to severe external barriers, a lack of market awareness, and fragile internal governance. Ultimately, a startup likelihood of survival depends on its *absorptive capacity*, specifically, its speed in capturing knowledge and accumulating industry-specific experience. This experiential learning improves the

firm's market positioning, mitigates external barriers, and strengthens internal governance, thereby securing its survival.

There is no universal agreement on the timeframe that defines startup survival. However, a widely accepted metric is the firm's ability to sustain payroll for its employees and owners for at least 42 months. A highly standardised operational model is provided by the Global Entrepreneurship Monitor (GEM), which sets out precise metrics to map the startup lifecycle:

- Birthday Event: reached by any enterprise that has paid its employees, or owners, for a minimum of three months.
- New Businesses: ventures that have maintained an active payroll for more than 3 months but less than 42 months.
- Actual Survival: achieved once a firm surpasses 42 months of continuous payments.

Beyond the GEM definition, researchers use various time horizons to study survival, depending on the sector and the study sample:

- Less than 3 years
- At least 3 years
- At least 4 years
- At least 5 years
- At least 6 years
- At least 7 years
- 8 or more years

Alongside these temporal definitions, the literature also offers conceptual alternatives to a purely time-based view of survival. Some studies treat survival as a multidimensional construct, shaped by four conditions: i. reaching the break-even point; ii. achieving accelerated growth; iii. maintaining a cash stock; iv. sustaining continuous operation. Josefy et al. (2017) propose a related but distinct framework. They break survival down into three analytical dimensions:

- **Operations:** the continuity of industrial or commercial activity.
- **Ownership:** the retention, or structured transfer, of proprietary assets.
- **Solvency:** the financial capacity to meet obligations.

Empirical work shows that survival is closely tied to founders' strategic decisions. When founders meet, or fail to meet, their own benchmarks (aspiration points and survival reference points), this shapes their choice to exit the market, maintain stability, or increase investment.

Given this heterogeneity of thresholds and definitions, this dissertation does not adopt a single, fixed survival cut-off. Instead, firm age is treated as a continuous duration variable. The hazard of exit will be considered at each point along this duration, and estimated econometrically through the Cox specification introduced in (see Chapter 4, Section 2.4).

This continuous-time approach is preferred over discrete thresholds, such as the GEM's 42-month rule or the shorter cut-offs reported in the literature (three to five years), for two main reasons. First, a fixed threshold treats survival as binary: achieved or not achieved at one point in time. This discards information on how the risk of exit changes across a venture's lifecycle. This information is central to the theoretical framework developed in this chapter. Second, a fixed threshold works poorly with staggered founding cohorts. Ventures founded more recently have had less time at risk of exit, simply by construction. A common cut-off would therefore confuse genuine survival differences with mechanical censoring.

This approach is also compatible with the conceptual dimensions outlined above. It does not force survival into a single measure. Instead, it allows operational cessation and ownership transition through acquisition to be modelled as distinct, cause-specific events. This is consistent with Josefy et al. (2017) distinction between operations and ownership, and it is developed further in the competing-risks framework introduced in Section 2.3.

2.1.3 The Valley of Death

In high-technology and highly innovative sectors, the early-stage survival of a venture is intrinsically synonymous with navigating the "Valley of Death". This critical phase refers to the highly vulnerable developmental and funding gap that elapses between the generation of basic scientific knowledge or R&D and the successful commercialisation of marketable goods or services. Empirically, this transitional corridor typically spans the first five years of operation, during which startups exhibit a severe lack of liquidity (OECD, 2015; Alunni, 2020). This is primarily because they are unable to generate sufficient revenues from commercial operations to offset their highly intensive R&D and initial operational costs.

The structural composition of a startup's founding team represents a critical determinant in navigating the financial and operational bottlenecks of the "Valley of Death". Initial resource scarcity often triggers an organisational "*snowball effect*", severely disrupting internal operations, structural development, and human resource acquisition. Under these constraints, solo-founded ventures are highly susceptible to mismanagement and technical lock-in, particularly where the founder's expertise is concentrated in a narrow technical domain, leaving commercial, financial, or business development functions under-resourced. Conversely, co-founding teams leverage a strategic division of labour and complementary capabilities across technical, administrative, and strategic domains. This structural synergy not only mitigates the risk of "entrepreneur discouragement" during liquidity crises but also enables the venture to articulate a 'credible business proposition', which is a recognised prerequisite for securing essential external backing from venture capitalists and business angels.

This organisational synergy is deeply tied to the founders' specific human capital. While strong technical qualifications are vital to sustain R&D-intensive product development, a lack of financial and managerial training often precipitates fatal cash-flow and accounting failures. This risk is moderated by the founders' professional background. Prior industry experience and pre-existing trust networks significantly facilitate early-stage investment decisions, while familiarity with regulatory or compliance requirements and complex sales cycles accelerates the path to commercialisation.

Additionally, the probability of a startup surviving the Valley of Death is highly sensitive to its initial size and resource scale. Small, highly resource-constrained ventures (e.g., micro-enterprises) are disproportionately vulnerable to the negative effects of financial constraints during their first year of operation. A smaller organisational size limits a venture's capacity to absorb the high fixed costs associated with R&D and product testing. Furthermore, smaller startups face more severe credit-rationing from traditional banking institutions, particularly where the underlying product is intangible in nature and therefore ill-suited to satisfying rigid collateral requirements, forcing founders to rely on more volatile personal or informal funding sources.

The lifecycle of startup debt (Hejazi and Seyyedamiri, 2024) demonstrates a complex, non-monotone relationship with survival:

- The startup stage: access to bank loans and external debt serves as a vital tool to foster growth and improve survival rates during the very early years.
- The post-entry stage (around the sixth year): the impact of debt obligations reverses. Accumulating excessive debt increases interest burdens, delays critical investments, and significantly increases the probability of shutdown and forced market exit.
- Loan source efficiency: interestingly, empirical models suggest that non-bank debt is twice as effective as traditional bank loans in predicting startup survival, primarily because most early-stage entrepreneurs face credit-rationing and cannot satisfy the strict collateral requirements of traditional banking institutions.

Finally, the literature suggests that crossing the Valley of Death requires active participation from specialized supporting actors. Rather than relying on traditional business incubators that merely provide generic office space and basic advisory services, high-tech ventures require active acceleration and incubation models that participate directly in profit generation and value co-creation. These models, coupled with Corporate Venture Capital (CVC) and Business Angel Networks, provide the crucial combination of smart money, regulatory navigation support, and market validation necessary to sustain the startup until it achieves commercial scale.

2.2 The relationship between founder background and vc capital investment

2.2.1. The high-opportunity-cost founders as a double- edged sword of capital

In the lifecycle of entrepreneurial ventures, the matching process between the founder's background and the acquisition of venture capital (VC) requires a deep strategic alignment of incentives, risk tolerances, and control structures.

A critical dimension of founders' background is their accumulated human capital (measured by pre-entry work experience and educational credentials), which directly shapes their entrepreneurial opportunity cost. Founders with highly accomplished backgrounds, such as those coming from established corporate tracks, forgo lucrative external career options to launch a startup (Arora, Nandkumar 2009). This high opportunity cost fundamentally alters their strategic decision-making and financing requirements:

- Strategic impatience and the demand for capital: high-opportunity-cost founders exhibit a form of "strategic impatience". Because lingering in a stagnant or moderately profitable venture is highly costly to them, they strategically reject the "steady state" approach of lifestyle businesses. Instead, they focus aggressively on achieving a rapid and highly lucrative liquidity event, a "cash-out" via an Initial Public Offering (IPO) or a high-value acquisition.
- The aggressive investment pathway (Burn Rate): to compress the timeline to exit, these experienced founders implement aggressive growth strategies characterized by high initial burn rates. This aggressive posture demands significant upfront capital injections, driving a high demand for external venture capital.

This experience-driven strategy is highly aligned with the investment horizon of VCs, who also seek high-velocity, high-upside exits. Empirically, this aggressive investment strategy financed by VCs leads to a distinct polarization of outcomes: it significantly accelerates the *hazard rate of successful cash-outs* but simultaneously

increases the risk of rapid "flame-out" (shutdown) due to premature scaling or strategic errors.

2.2.2 Macroeconomic aspects: hot and cold markets

The strategic matching process between founder backgrounds and venture capital investment is fundamentally moderated by the macroeconomic environment and cyclical shifts in capital supply. The criteria used by financial intermediaries to evaluate entrepreneurial human capital are not static. They change systematically as capital abundance alters the market's underlying incentives, with consequences that extend beyond the initial financing decision to the venture's eventual exit trajectory.

During venture capital peaks, or "hot markets", driven primarily by positive supply-side shocks of capital from institutional limited partners, the marginal cost of experimentation for early-stage investors falls significantly. This capital influx relaxes financial constraints and shifts VCs' focus on radical innovation. VCs behave as rational experimenters, showing a pronounced willingness to overlook conventional background constraints. During these expansionary waves, they redirect funding toward riskier, more technologically novel ventures led by younger, less-proven founders, and they rely on smaller, more concentrated syndicates as the wide availability of capital reduces financing risk. This cyclical shift in risk tolerance is particularly marked among the most experienced venture capitalists, who actively pivot toward high-variance, novel investments when capital is abundant.

This same influx of capital, once deployed, also shapes the venture's ultimate trajectory, producing a pronounced "twist" in the distribution of exit outcomes (Nanda and Rhodes-Kropf, 2013). In hot markets, VCs' propensity to back disruptive, ex-ante risky technology translates into substantially higher ex-post failure and shutdown rates. Paradoxically, however, the same capital abundance also engineers the most extraordinary successes. Startups that survive and exit during hot markets achieve far higher valuations, whether via IPO or acquisition,

and leave a stronger innovation footprint, reflected in higher patent filings and subsequent citations.

Conversely, during contractionary "cold markets", the cost of failure rises, and financing risk intensifies. VCs demand dual risk insurance: they revert to "safer", highly mature founders with extensive industry backgrounds, and they construct larger, more diversified syndicates to spread the financial burden. Ventures funded in cold markets show a higher baseline survival rate, but they struggle to secure multi-billion-dollar exits or high-premium acquisitions, typically following more restrained, linear growth paths instead.

Thus, the macroeconomic cycle acts as a critical institutional moderator at both stages of the venture's lifecycle. It determines whether VCs demand highly experienced founders or bypass traditional human capital requirements to fund radical innovation (Nanda and Rhodes-Kropf, 2013), and it subsequently amplifies outcome variance, driving hot-market ventures toward the extreme tails, either exceptional success or outright failure, of the ex-post-performance distribution. Securing substantial equity backing during an economic expansion does not, in itself, ensure a firm's survival. Rather, it raises the stakes in both directions.

Finally, the founders' background in terms of their initial financing decisions dictates whether they can attract VC equity later in their lifecycle. A founder's early choice to secure long-term debt financing from outside creditors (such as banks) rather than equity creates persistent path-dependencies:

- The debt-constrained cap table: startups that rely heavily on initial long-term bank debt face strict ongoing repayment and interest obligations, reducing operational flexibility (Honjo, Iwaki, Kato, 2024).
- Private Equity aversion to debt: more importantly, VCs and private equity investors are highly averse to investing in or acquiring startups with significant pre-existing bank debt. This is because outside lenders (such as main banks) hold senior claims and established credit relationships that directly conflict with the VC's demand for exclusive corporate control. Consequently, a founder's decision to use bank debt at founding effectively

forecloses their ability to raise subsequent venture capital or achieve a successful exit via merger (Honjo, Iwaki, Kato, 2024).

Having established how founder background and the prevailing capital cycle jointly determine a venture's capacity to attract VC funding, and how that same cycle shapes the venture's eventual exit, the following section turns from financing to exit itself: specifically, how the presence and magnitude of VC investment determine the mode and timing of a venture's ultimate market outcome.

2.3 The relationship between the VC capital investment and the mode of exit

In the traditional entrepreneurship theories, firm survival has been regarded as the paramount indicator of positive performance failure (Arora, Nandkumar 2009). On the other hand, contemporary paradigms within corporate finance and entrepreneurial strategy emphasises a high heterogeneity in new venture exit pathways, drawing sharp distinctions between insolvency-driven cessation (shutdown), voluntary liquidation, Initial Public Offering (IPOs), and Mergers and Acquisitions (M&A). Within this theoretical framework, the acquisition of substantial external financial resources does not merely foster firm survival. Rather, it fundamentally alters the founder's strategic horizon, accelerating the venture's trajectory toward extreme and polarized outcomes (Honjo et al., 2024). Of these four pathways, this dissertation focuses empirically on the two that are both economically distinct and reliably observable within the sample: shutdown and acquisition. Voluntary liquidation and IPO are considered separately below.

2.3.1 “Cash-Out” and “Flame-Out”

The infusion of venture capital fundamentally reorients the strategic horizon of a new venture. Arora, Nandkumar's study (2009) argues that in innovation-based industries, prolonged survival is rarely an end in itself. Rather, growth-oriented founders and their VC backers prioritise achieving a substantial liquidity event, or 'cash-out' (via an IPO or a favourable acquisition). Venture capitalists provide the risk capital necessary to sustain an aggressive “burn rate”, which accelerates product development and market footprint. Although this VC-backed strategy

greatly improves the chances of a successful cash-out, it operates as a double-edged sword.

Attempting to achieve a rapid exit while spending cash rapidly leaves almost no margin for operational error. This high-pressure environment increases the risk of complete failure, or “flame-out”, often caused by premature scaling or product design errors. As a result, VC investments lead to highly polarised, extreme outcomes. The startup either achieves a major success (IPO or strategic acquisition) or fails completely, rarely surviving as a slow-growing, stable business (Arora, Nandkumar 2009).

This drive towards high-value exit routes, particularly mergers and acquisitions, is deeply consistent with the VC-driven growth mandate and shapes the earliest positioning decisions of the firm. Scalable, growth-oriented ventures deliberately select specific financial and strategic positioning paths to make themselves highly attractive to external equity investors. A prominent example of this ex-ante signaling is the choice of naming strategy: scalable entrepreneurs systematically eschew eponymous naming conventions (i.e. naming the firm after the founder's surname) (Belenzon et al., 2020). By avoiding eponymous branding, founders signal organizational scalability to financiers, trading off the immediate reputational and profitability benefits of personal 'glory' in favour of attracting outside capital for “growth” (Belenzon et al., 2020). Empirically, these non-eponymous, highly scalable firms grow twice as fast as their eponymous counterparts, experiencing average sales growth of 4.1% compared to 2.1%, and demonstrate a significantly greater propensity to experience successful liquidity events, such as major acquisitions or IPOs (Belenzon et al., 2020). Thus, the strategic anticipation of attracting venture capital influences not only the ex-post survival or failure rates of the firm but also dictates ex-ante structural choices designed to channel the venture towards an M&A or IPO exit route (Belenzon et al., 2020).

2.3.2 The role of macroeconomic cycles

As argued in Section 2.2.2, the macroeconomic capital cycle does not only shape a venture's initial access to financing, it also moderates the value and mode of its eventual exit, producing the polarised, cohort-dependent outcomes discussed

above. This dynamic underpins the cohort fixed effects introduced in the empirical strategy in Section 2.4.

2.4 Summary and theoretical framework

The literature reviewed in this chapter converges upon a single, sequential narrative describing the trajectory of a new venture, articulated across three interconnected theoretical strands.

The first strand (Section 2.1) argued that **organisational survival** is not a binary or static condition, but a multidimensional, time-dependent construct shaped by the firm's stage of development. Ventures navigating the earliest phases of their lifecycle, and in particular the "Valley of Death", face an elevated baseline hazard rate driven by the rapid obsolescence of founding endowments, resource scarcity, and financial constraints. Crucially, this literature identified the founding team's structural composition, human capital, and access to external capital as the principal levers through which this early hazard can be mitigated.

The second strand (Section 2.2) argued that **founder background** does not operate as an isolated determinant of survival, but as a strategic signal that shapes the venture's capacity to attract venture capital. High-opportunity-cost founders pursue aggressive, capital-intensive growth strategies aligned with VC return expectations, whilst the macroeconomic capital cycle systematically moderates the threshold of founder experience that investors are willing to accept. Founder background, in this sense, functions as a mediating variable between the firm's initial endowments and its subsequent access to equity financing.

The third strand (Section 2.3) argued that the acquisition of VC capital is not merely a facilitator of survival, but an accelerant that polarises the venture's ultimate trajectory towards extreme outcomes, namely a successful "cash-out" (IPO or acquisition) or an outright "flame-out" (shutdown). This polarising effect is, once again, moderated by the prevailing macroeconomic cycle, which alters both the probability and the magnitude of the eventual exit.

Taken together, these three strands describe a **sequence** running from the venture's initial conditions, through the founder's strategic positioning and financing decisions, to its ultimate exit trajectory, with survival acting as the necessary precondition that must hold at each stage for the chain to proceed.

CHAPTER 3: THE CONTEXT OF ANALYSIS: THE CYBERSECURITY INDUSTRY

3.1 Introduction

The theoretical framework outlined in Chapter 2 was deliberately formulated at a level of generality applicable to high-technology ventures. The present chapter narrows this analytical lens to focus specifically on the cybersecurity sector, examining how the general mechanisms of survival, founder-background signalling, and VC-driven exit polarisation are manifested within a context, the cybersecurity industry, characterised by distinctive financing cycles, acquirer structures, and innovation dynamics. This sector-specific contextualisation constitutes a necessary precondition for the empirical operationalisation adopted in the analysis that follows.

3.2 Market size and growth trends

The global cybersecurity market has expanded substantially over the past decade. Industry estimates place the market at approximately USD 75 billion in 2015, with investment in the sector growing from approximately USD 1.1 billion to USD 3.8 billion between 2011 and 2015 (Zhang, 2016, citing Press, 2015 and CBInsights, 2016). Post-pandemic industry projections estimate a continued expansion from USD 217.9 billion in 2021 to USD 345.4 billion by 2026, corresponding to a compound annual growth rate of 9.7 % (Moiseienko and Kiva, 2022). Longer-term forecasts cited within the industry anticipate global security expenditure surpassing USD 300 billion by 2025.

This sustained growth is closely tied to the expansion of the Internet of Things and the corresponding growth of the “global attack surface”. The number of connected devices worldwide stood at approximately 10.94 billion in 2020 and is projected to reach 25.44 billion by 2030 (Moiseienko and Kiva, 2022). These trends underscore the sector's structural growth momentum and are consistent with the elevated financing propensity observable within more recent cohorts in the empirical sample.

3.3 Structural characteristics of the cybersecurity startup ecosystem

3.3.1 Market fragmentation

To frame the sector under analysis, it is useful to first outline its competitive structure. Industry literature almost unanimously describes this structure as highly fragmented. No single security platform is universally applicable. The market instead consists of a multitude of startups, each specialised in narrow “attack vectors” that define specific market segments (Ramsinghani, 2016, as cited in Zhang, 2016).

This segmentation can be traced along three overlapping dimensions identified by analysts: the type of security addressed, the solution category offered, and the type of service delivered. These dimensions range from network, endpoint and cloud security, to encryption, threat management, and managed security services. Zhang's (2016) synthesis of *Markets and Markets* data shows that this taxonomy is not purely descriptive. It has concrete operational implications: firms combining multiple sub-categories in their offering face increasing layers of complexity, since a single solution area (such as data loss prevention) can itself be delivered through several distinct service models (design and integration, consulting, and so on).

The granularity of this classification becomes even more evident at the sub-segment level. Within Identity and Access Management alone, the category breaks down into six further components: provisioning, directory technologies, single sign-on, password management, advanced authentication, and audit/compliance/governance. Each of these is served by a distinct cluster of players (solution developers, resellers, IAM-as-a-service providers, and professional service firms). This pattern of "nested" fragmentation is broadly consistent across other cybersecurity sub-sectors. It also helps explain why even large, diversified vendors rarely operate as pure-play specialists: firms such as Microsoft compete simultaneously across multiple, unrelated security categories, further blurring segment boundaries.

Table 3.1: market structure

Dimension	Segmentation (level 1)	Sub-segmentation (level 2)
The type of security addressed	Network security	—
	Endpoint security	—
	Cloud security	—
	Encryption	—
The solution category offered	Threat management	—
	Identity and Access Management (IAM)	Provisioning
		Directory technologies
		Single sign-on (SSO)
		Password management
		Advanced authentication
		Audit / Compliance / Governance
	Data Loss Prevention (DLP)	—
The type of service delivered	Design and integration	—
	Consulting	—
	Managed security services	—

It is precisely this evidence of fragmentation, and the resulting difficulty in assigning sample firms to mutually exclusive, theoretically grounded categories, that motivates the clustering approach adopted in this study. Firms in the sample were grouped into a limited number of internally coherent macro-categories.

This methodological choice reflects a trade-off identified in the literature. On one hand, a very detailed segment variable (for example, splitting solutions into categories such as encryption, SIEM, or endpoint) would capture competitive

differences more precisely. However, this risk creating cells with very few observations, since most segment-specific startup populations are small. Zhang's (2016) own data illustrate this problem. Even at the aggregate industry level, the number of M&A and IPO transactions is modest.

On the other hand, a very broad categorisation risks flattening important competitive distinctions. The fragmentation literature points, for instance, to a meaningful divide between horizontal/platform players (broad, multi-segment providers, more likely to grow organically or pursue an IPO) and vertical/niche players (specialists in a single attack vector, more often acquired through M&A by larger vendors seeking complementary capabilities).

The clustering process adopted in this thesis therefore aimed for an intermediate compromise. Macro-clusters were defined to be thematically coherent, while remaining broad enough to ensure adequate statistical size within each category. The final cluster in this scheme serves a residual function. It was constructed to capture highly specialised sector verticals, digital services peripheral to the cybersecurity core, and early-stage entities whose structural characteristics deviate from the mainstream markets identified in the main clusters.

The Residual Category was constructed to preserve the thematic integrity of the three primary macro-sectors (Core Cybersecurity & Risk Management, Software & Deep Tech, and IT Services & Professional Consulting) by isolating the three sub-categories listed above, which fall outside their core scope.

- *Fintech & Insurtech*: disruptive technology providers operating in the financial and insurance verticals, with a focus on digital payment rails, neo-banking, and automated actuarial modelling.
- *Digital Media, E-Commerce & Agency Services*: firms specialised in digital marketing, online retail marketplaces, and creative content management, driven by data-centric user acquisition strategies.
- *Startup & Research Entities*: early-stage ventures, academic spin-offs, and research-oriented organisations, characterised by high experimental risk,

pre-revenue operations, and a structural focus on long-term technological breakthroughs rather than immediate market scaling.

This clustering structure does not fully resolve the tension between granularity and statistical robustness highlighted in the literature. It does, however, preserve enough information in the sector-segment control variable to distinguish the main competitive logics of the industry, while avoiding the sparsity problems typical of excessive disaggregation.

3.3.2 The "cockroach" versus the "unicorn" survival model

Cybersecurity startups also differ from other technology verticals in their underlying growth narrative. In many segments of the tech industry, entrepreneurial success is closely linked to the pursuit of rapid "unicorn" status, that is, a valuation reaching USD 1 billion. This narrative rewards aggressive top-line growth and rapid user acquisition. It also rewards a willingness to sustain significant losses in the pursuit of market share, typically financed through successive rounds of venture capital.

Cybersecurity startups follow a different pattern. Within the venture capital community, this alternative growth archetype is informally known as the "cockroach" model (Ramsinghani, 2016 as cited in Zhang, 2016). Cockroach companies do not prioritise speed of expansion. Instead, they are defined by their capacity to survive adverse conditions, such as funding downturns, market volatility, and prolonged periods without external capital, while maintaining a path towards profitability or, at minimum, capital efficiency.

This distinction reflects structural features of the cybersecurity market itself. Demand for security solutions tends to be relatively inelastic and recurring. It is driven by regulatory compliance requirements, the continuous evolution of threat vectors, and the reputational and financial costs associated with breaches. As a result, cybersecurity firms are often able to generate stable, subscription-based revenue streams. This holds even in the absence of the explosive user growth that

typically characterises consumer-facing technology unicorns (Ramsinghani, 2016, as cited in Zhang, 2016).

This structural resilience is broadly consistent with the comparatively low baseline shutdown incidence observed empirically within the sector. It also offers a substantive, sector-specific explanation for one of the patterns that will emerge from the analysis in this thesis. Financing propensity does not, on its own, appear to function as a straightforward protective or penalising factor against outright firm failure. If cybersecurity firms are structurally predisposed towards survival, through capital efficiency and operational flexibility rather than through capital intensity alone, then the amount of external financing raised may carry less explanatory weight for survival outcomes than it typically does in other, less resilient technology sectors.

3.3.3 Venture Capital role and mode of exit

Organisational survival in high-technology entrepreneurship cannot be treated as a simple binary outcome. Venture-backed startups face highly heterogeneous exit pathways. In the cybersecurity industry, this divergence is best captured through a **competing-risks survival framework**. This framework distinguishes between two mutually exclusive, cause-specific events: **strategic acquisition**, representing a successful, VC-aligned "cash-out", and **operational bankruptcy or shutdown**, representing an insolvent "flame-out". Public listings (IPOs) are excluded from this taxonomy as they remain extremely rare in the cybersecurity sector. Acquisition therefore stands as the primary observable proxy for strategic success.

Empirical evidence, that will be presented in the next chapter, confirms that exits in this sector occur predominantly through mergers and acquisitions. Two distinct classes of strategic buyers drive this pattern. The first are **"OS-layer" acquirers**: large technology platform providers such as Microsoft, Intel, and Google. These firms typically purchase niche startups to integrate specialised security capabilities directly into core operating systems and enterprise platforms. Acquisition-volume data compiled by AGC Partners identifies IBM, Intel, and Microsoft as the most active corporate buyers by transaction count between 2008 and 2016. The second

category comprises "**carriers**": telecommunications operators such as AT&T and SingTel. These firms acquire cybersecurity companies to upsell integrated security services alongside their managed cloud and data-centre infrastructure.

The presence of these acquirers, together with active Corporate Venture Capital (CVC) arms, gives cybersecurity startups a sector-specific route to strategic exit. This route helps them navigating the financial bottlenecks commonly referred to as the "Valley of Death."

Exit horizons, however, remain structurally bimodal. Smaller, niche point solutions tend to exit relatively quickly, often within the first few years of founding, while ventures that scale towards very large valuations typically require a much longer gestation period, unfolding over the better part of a decade or more. This divergence between rapid, smaller-scale exits and long-gestation, high-value strategic acquisitions is a recurring pattern in the sector's M&A activity.

There are some other characteristics of the cybersecurity industry that motivate some of the choices that will be made in the empirical analysis in the following chapters. First, this bimodal pattern of exit provides the theoretical basis for controlling for **cohort fixed effects** in the Cox estimations. A startup's survival and exit trajectory are shaped by the specific "vintage" of venture capital available at founding. During expansionary "**hot markets**", capital is abundant. VCs behave as active experimenters. They fund riskier, unproven founding teams. This simultaneously raises the hazard of a high-premium strategic exit and the hazard of rapid shutdown, driven by high cash-burn rates and premature scaling. Cohorts founded in "**cold markets**" face tighter underwriting standards. These cohorts show higher baseline survival, but a lower probability of reaching multi-billion-dollar exits.

Cohort fixed effects also capture ex-ante regulatory shocks. These shocks alter market-wide acquisition propensities independently of firm-level performance. A clear sector-specific example is the introduction of the North American Electric Reliability Corporation's Critical Infrastructure Protection standard, version 5 (NERC-CIP v5), announced in late 2013. This regulatory announcement triggered

a temporary, highly concentrated surge in M&A activity. Large-regulated entities rapidly acquired compliant cybersecurity startups to meet compliance deadlines, rather than developing internal capabilities.

3.3.4 The nature of innovation

According to Tushman and Anderson (1986), industries evolve through long periods of incremental improvement, interrupted by rare but decisive technological breakthroughs. These breakthroughs, or discontinuities, sharply alter existing price-performance relationships. Crucially, not all discontinuities have the same effect on incumbent firms. Some build on existing knowledge. Others render that knowledge obsolete.

Competence-enhancing discontinuities represent significant improvements that build on know-how already present within an industry. They substitute for older technologies without displacing the skills required to master them. As a result, established firms are typically best placed to exploit these advances. This tends to consolidate their market position, raise barriers to entry, and reduce turbulence within the sector.

Competence-destroying discontinuities operate in the opposite direction. They require fundamentally new skills, knowledge, and operating routines. Mastery of the new technology bears little relation to mastery of the old one. For this reason, such discontinuities are typically initiated by new entrants rather than incumbents. Established firms, constrained by sunk costs, organisational routines, and prior technological commitments, often struggle to adapt. This pattern favours new entrants at the expense of firms tied to legacy technology, and it increases both competitive turbulence and the hazard of exit for incumbents.

This distinction maps closely onto developments within the cybersecurity sector. Early cybersecurity products were built around signature-based detection: firewalls and antivirus software that identified threats by matching them against known patterns. Improvements to this paradigm, faster signature updates, broader threat libraries, better integration with existing IT infrastructure, are properly understood

as competence-enhancing. They extended an established technological logic without displacing it. Established vendors, already proficient in this logic, were well placed to lead such incremental advances.

The shift towards data-driven security engineering represents a discontinuity of a different kind. Cloud-native security architectures, behavioural and anomaly-based threat detection, big data analytics, and machine-learning-driven forensics do not simply extend signature-based logic. They replace it with a different technical and conceptual foundation. Mastering these approaches requires skills in distributed systems, statistical modelling, and applied machine learning, competences that differ substantially from those underlying legacy security products. Consistent with Tushman and Anderson's argument, this transition has disproportionately favoured new entrants, while exposing incumbents tied to older architectures to heightened competitive pressure.

This theoretical distinction will underpin the inclusion of **sector-segment fixed effects** in the empirical model. Firms operating within data-driven, competence-destroying segments of the market are expected to face a structurally different hazard profile from firms operating within more consolidated, competence-enhancing segments. Without this control, segment-level technological turbulence could be mistaken for firm-level weakness in governance or founder human capital, undermining the interpretation of the survival estimates presented in Chapter 4.

3.4 The Valley of Death in cybersecurity ventures

Within cybersecurity, the vulnerability associated with the Valley of Death, discussed in general terms in Section 2.1.3, is especially severe. The sector is marked by rapid technological obsolescence, so R&D cycles are both capital-intensive and subject to high technical uncertainty. The main cause of failure at this stage is the acute shortage of risk investment. Institutional venture capitalists and traditional financial institutions are often reluctant to invest during the early, uncertain stages of a venture, given the combined burden of technical and commercial uncertainty. As a result, a cybersecurity startup's chances of survival depend heavily on its financial structure and its ability to secure external capital

over time. Strong initial funding at founding is a good predictor of long-term viability. However, financial constraints and cash-flow disruptions during the first year of operation have the most severe negative effect on a firm's probability of survival (Zapata-Molina and Montes-Hincapié, 2022).

This vulnerability is made worse by dynamics within founding teams. As noted in Section 2.1.3, solo-founded ventures are generally at high risk of mismanagement and technical lock-in when the founder's expertise is limited to a narrow technical area. This risk is particularly strong in cybersecurity, where solo founders are mostly technical specialists, such as software engineers or cryptographers, rather than trained commercial operators. The lack of complementary commercial, financial, or business-development skills increases commercial uncertainty and raises the risk of a mismatch between product and market. This supports the theoretical argument that founding-team composition should be treated as a determinant of survival, not merely a secondary characteristic.

The severity of credit rationing, discussed in general terms in Section 2.1.3, is also greater in this sector. Because cybersecurity products are, almost by definition, intangible, software- or service-based assets, small cybersecurity startups are structurally unable to offer the physical collateral required by traditional banks. This forces early-stage founders to rely heavily on volatile personal or informal funding sources, reinforcing the sector's dependence on specialised early-stage risk capital rather than debt financing.

3.5 Overview and final discussion

The evidence reviewed in this chapter has shown how cybersecurity is an heterogeneous and structurally uneven industry. The industry combines rapid overall growth with wide variation across segments, in deal activity, in investment size, and in the type of firm that attracts capital. This fragmentation reflects genuine differences in competitive positioning between broad, multi-segment providers and narrow, single-purpose specialists, and it shapes which firms are able to attract external financing in the first place.

This sector-specific picture builds upon the general theoretical framework outlined in Chapter 2, where the literature described a sequential chain running from a

venture's initial conditions, through founder characteristics and financing decisions, to its eventual market exit, with survival as the condition that must hold at each stage. The discussion in this chapter aligns with this picture. Capital is heavily concentrated in a small number of segments, cybercrime costs and demand pressures drive sharp swings in VC funding, and the divide between specialist and generalist firms offers a plausible, sector-specific explanation for the polarised outcomes, acquisition versus shutdown, discussed in Section 2.3.

Taken together, the theoretical framework discussed in Chapter 2 and the characteristics of the industry reviewed in this chapter lead to the formulation of the following research questions of this thesis:

1. *What variables influence the likelihood of receiving funding for cybersecurity startups?*
2. *Conditional on being funded, what are the determinants of survival of cybersecurity startups?*
3. *Conditional on being funded, in what way firms exit the market, (whether through, acquisition, or failure)?*

The following chapter operationalises these questions. It sets out the variables used to capture founder characteristics, financing status, and sector structure, and specifies the two-stage empirical strategy, financing propensity (Logit), survival and hazard estimation (Cox), and the competing-risks extension, through which the research questions will be addressed.

CHAPTER 4: DATA AND METHOD

4.1 Introduction and general description of the dataset

The empirical analysis presented in this thesis draws upon a dataset constructed through a targeted extraction of firm-level data from Crunchbase. The extraction was conducted using the set of category tags and keywords outlined in the preceding section, which were selected to capture the multidimensional and fragmented nature of the cybersecurity sector.

Six thematic categories were employed to guide the search strategy.

- I. The first, *Core FinTech*, encompassed firms tagged under Financial Services, Digital Banking, Payments, and Lending Platforms, capturing core financial innovation and digital entrants operating within traditional banking functions.
- II. The second, *RegTech*, comprised firms classified under Regulatory Technology, Compliance Software, and Risk Management, reflecting the development of encryption-based, data protection, and regulatory compliance solutions.
- III. The third, *InsurTech*, included firms tagged as Insurance Technology and Digital Insurance Platforms, illustrating the diversification and vertical fragmentation characteristic of the financial sector.
- IV. The fourth category, *Cybersecurity and Encryption*, formed the analytical core of the extraction, drawing on the tags Cybersecurity, Data Security, Encryption Software, and Privacy Tech, directly addressing the adoption of encryption as a strategic asset.
- V. The fifth, *Blockchain and Crypto Infrastructure*, covered firms under Blockchain, Crypto Infrastructure, and Digital Assets Security, representing new entrants driving technological transformation within finance.
- VI. The sixth, *WealthTech and Robo-Advisory*, encompassed firms tagged under Wealth Management Technology and Investment Platforms, illustrating the digital variety and specialised niches emerging within the sector.

In addition to these thematic categories, four complementary set of keywords were applied to further refine the sample further and to strengthen its relevance to the specific focus of the thesis. These keywords comprised: RegTech and Compliance (Regulatory Technology, Compliance Software, Risk Management), capturing small and medium-sized enterprises offering encryption-driven solutions; FinTech Infrastructure, encompassing non-consumer-facing firms tagged under Financial Software, Payments Infrastructure, and Banking Software, representing business-to-business entities operating in partnership with banking institutions; IT Services for Finance, comprising firms tagged as Enterprise Software, IT Services, and Cloud Security, a category under which a considerable share of Italian SMEs are classified; and, where available, Digital Trust Technologies, including Identity Management, Authentication, and Secure Communications tags, directly linked to the themes of trust and firm survival explored in this study.

The observation window for the dataset spans the period *2009–2024*, a timeframe that captures both the early emergence of the fintech phenomenon in the aftermath of the global financial crisis and its subsequent maturation through successive waves of digitalisation, regulatory adaptation and, more recently, blockchain-based innovation. The geographical scope of the analysis is restricted to firms headquartered in *Europe and the United Kingdom*.

The initial extraction yielded a gross sample of firms matching the selected keyword categories, which was subsequently subjected to a data-cleaning process aimed at ensuring the internal consistency and relevance of the sample to the research objectives. Specifically, the following exclusion criteria were applied:

- I. Firms whose core activity was predominantly oriented towards marketing-related services, rather than financial technology *per se*, were excluded, as their inclusion would have diluted the sectoral focus of the analysis.
- II. Firms identified as short-lived spin-offs, defined as ventures with an operative history of six months or less, were excluded on the grounds that such entities do not provide a meaningful basis for the analysis of firm-level trajectories, founder characteristics or financing patterns.

Following the application of these criteria, the final sample comprises *1,227 firms*, constituting the empirical basis for the descriptive and econometric analysis presented in the remainder of this thesis.

The unit of analysis is the individual firm, observed at a single point in time. The baseline is therefore cross-sectional rather than longitudinal, with firm-level attributes recorded as of the most recent data extraction. Founder-related variables require a finer level of detail. For these, individual founders within each firm are examined to describe team composition and professional and educational background. These founder-level observations are then aggregated to the firm level, so that the descriptive statistics and the regression analysis are conducted at firm level.

The dataset includes a broad set of variables:

- I. firm demographic and operational characteristics, including revenue and employee-size classes, operating status (i.e. active/ non-active etc);
- II. founder-level characteristics, including: team composition and professional and educational background;
- III. financing history, including funding round typology and count;
- IV. financial outcome variables, most notably total equity raised, which constitutes the dependent variable in the logit regression presented in the following chapter.

A full overview of the variables included in the dataset, together with their type and source, is shown in Table 4.1.

Table 4.1: Overview of dataset variables


Variable	Type	Description
<i>Cluster_idx</i>	4 categories	Core cybersecurity & risk management, IT services & professional consulting, software & deep tech, residual category
<i>Revenue class</i>	5 classes	Less than \$1M, \$1M to \$10M, \$10M to \$100M, \$100M to \$500M, \$500M+
<i>Employee class</i>	4 classes	0-10, 11-200, 201-1000, 1001+
<i>Operating status</i>	Dummy	Open = 1; Closed = 0
<i>Closure type</i>	Dummy (conditional on closed)	Acquired = 1; Shutdown = 0
<i>Number of founders</i>	Count	Solo founder vs team
<i>Founder professional background</i>	Dummy	Relevant = 1 if prior employment was in the same sector; 0 otherwise
<i>Founder educational background</i>	Dummy	Relevant = 1 if academic training was in a related field; 0 otherwise
<i>IPO status</i>	Dummy	Public = 1; Private = 0

Variable	Type	Description
<i>Investment stage</i>	Categorical	Bootstrapped, Early-stage, Growth-stage, late stage
<i>Number of funding rounds</i>	Count	—
<i>Total equity amount</i>	dummy	1 = financed; 0 = bootstrapped
<i>cohorts</i>	4 categories	2009-2012, 2013-2017, 2018-2022, 2023-2026

4.2 Cluster variables

As mentioned in Section 3.3.1, the firms in the dataset were subsequently grouped into four macro-categories, or clusters, on the basis of their prevailing business model and core value proposition. This classification was undertaken in order to move beyond the granularity of the original Crunchbase category tags and to construct a more analytically tractable taxonomy suited to the objectives of the present study.

Figure 4.1: overview of companies' type



As illustrated in Figure 4.1, the resulting distribution across the sample of 1,227 startups is as follows: *Core Cybersecurity & Risk Management* constitutes the largest group, accounting for 390 firms; it is followed by the *Residual category*, comprising 335 firms; *Software & Deep Tech* accounts for 297 firms; and *IT Services & Professional Consulting* represents the smallest group, with 205 firms. This distribution indicates that firms whose core business is explicitly and narrowly focused on cybersecurity and risk mitigation represent the single most prominent group within the sample, while the remaining three clusters together account for a broadly comparable, if somewhat more fragmented, share of the dataset.

The identification of these four macro-categories followed a bottom-up, iterative process. No fixed taxonomy was imposed from the outset. Instead, the process began with the raw category tags and firm descriptions extracted from Crunchbase. These were first grouped into narrower sub-sectors, based on each firm's main business focus. The sub-sectors were then grouped into the four broader macro-categories. This grouping was based on their thematic closeness and strategic overlap. Firms with a very low frequency, or an unclear position, were placed in a

separate residual cluster. This helped preserve the internal coherence of the three main categories.

The *Core Cybersecurity & Risk Management* cluster is the largest group in the dataset. It brings together firms whose core business is to protect the integrity, availability, and confidentiality of digital assets. These firms tend to be product-centric. They have a highly specialised strategic focus on risk mitigation. This macro-category consolidates four sub-sectors: i) *cloud security, compliance and cybersecurity*, encompassing specialised providers engineering secure distributed environments and ensuring adherence to international regulatory frameworks; ii) *cybersecurity and technology solutions*, comprising vendors of dedicated security hardware and software ecosystems, including endpoint detection, perimeter defence, and automated threat response tools; iii) *data and threat intelligence*, referring to organisations focused on proactive, information-led defence through the analysis of global threat vectors, dark web monitoring, and data-centric protection strategies; iv) *cybersecurity and risk management*, denoting strategic consultancy practices specialising in enterprise risk assessment, corporate governance, and the development of holistic security architectures.

The *Software & Deep Tech* cluster aggregates companies positioned at the frontier of technological innovation, in which proprietary software development, advanced computing architecture, or emergent digital paradigms constitute the primary business driver, thereby bridging fundamental technical research with commercial software engineering. This category includes four domains: : i) *artificial intelligence, data and deep tech*, comprising entities focused on applied machine learning, predictive modelling, and fundamental research in deep technology, where data architectures serve as the primary catalyst for innovation; ii) *artificial intelligence, cloud and infrastructure*, encompassing providers of the underlying computational power, high-performance computing, and architectural frameworks necessary for model training, including cloud-native orchestration and specialised hardware provisioning; iii) *emergent paradigms*, referring to distributed ledger technology providers and cryptographic framework developers, exemplified by blockchain; iv) *software development and engineering*, denoting firms specialised

in custom software architecture, systems engineering, and the development of core platforms designed for third-party integration.

The *IT Services & Professional Consulting* cluster distinguishes service-oriented implementation and high-level strategic advisory from pure-play product engineering, encompassing firms for which technology serves either as an operational asset to be managed or as a strategic enabler of broader business transformation, rather than as the commercialised product itself. This macro-category includes four sub-sectors: i) *IT infrastructure and cloud services*, comprising managed service providers and firms handling virtualisation, hosting, and the ongoing maintenance of enterprise IT ecosystems; ii) *consulting and technology services*, referring to system integrators and technical consultants focused on the localised, practical deployment of complex third-party technological stacks; iii) *professional services and consulting*, denoting high-level corporate advisory firms providing business strategy, legal compliance, and management consulting; iv) *SaaS and B2B enterprise solutions*, comprising providers of scalable, cloud-based business applications designed to optimise corporate workflows (such as CRM, ERP, and productivity suites) through service-based models.

Finally, the *Residual category* captures highly specialised sector verticals, peripheral digital services, and early-stage entities whose structural characteristics deviate from the three established macro-sectors described above. This cluster identifies four sub-groups: i) *FinTech and InsurTech*, comprising disruptive technology providers operating within financial and insurance verticals, with a focus on digital payment rails, neo-banking, and automated actuarial modelling; ii) *digital media, e-commerce and agency services*, denoting entities specialising in digital marketing, online retail marketplaces, and creative content management driven by data-centric user acquisition strategies; iii) *start-up and research entities*, referring to early-stage ventures, academic spin-offs, and research-centric organisations characterised by high experimental risk, pre-revenue operations, and a structural focus on long-term technological breakthroughs rather than immediate

market scaling. This last category accounts for the remaining share of firms not readily assimilable within the three primary macro-sectors.

Table 4.2: market fragmentation

Macro-clusters	Sub-categories included
<i>Core Cybersecurity & Risk Management</i>	• Cloud security, compliance and cybersecurity • Cybersecurity and technology solutions • Data and threat intelligence
<i>Software & Deep Tech</i>	• Artificial intelligence, data and deep tech • Artificial intelligence, cloud and infrastructure • Emergent paradigms • Software development and engineering
<i>IT Services & Professional Consulting</i>	• IT infrastructure and cloud services • Consulting and technology services • Professional services and consulting • SaaS and B2B enterprise solutions
<i>Residual Category</i>	• FinTech and InsurTech • Digital media, e-commerce and agency services • Start-up and research entities


4.3 TEMPORAL DIMENSIONS OF THE DATASET

4.3.1 Foundation period

The founding period spans from 2009 to 2024. Figure 4.3 shows the distribution of firms by year of foundation. New firm creation was relatively modest and stable between 2009 and 2012, with between 45 and 55 firms founded each year. From 2013 onwards, foundation activity rose steadily, reaching a peak in 2017, when 126 firms were founded, the highest figure in the sample. After 2018, the number of new firms declined gradually year on year, falling to around 90 by 2021 and to roughly 50–60 in 2022–2024. This later decline should be interpreted with some

caution, as more recent cohorts may be under-represented due to reporting lag in the Crunchbase database.

Figure 4.2: Distribution of companies by foundation year



This growth pattern fits well with the European regulatory and threat landscape of the time. The NIS Directive was the first horizontal EU cybersecurity law. It entered into force in August 2016 (Zygierewicz, 2020). The General Data Protection Regulation came into full application in May 2018. It required organisations across the EU to protect personal data against unauthorised access and breaches.

Regulation kept intensifying after this point. A wider cybersecurity package was presented in September 2017. The EU Cybersecurity Act then entered into force in 2019. It gave ENISA a permanent mandate and set up the EU's first cybersecurity certification framework (Zygierewicz, 2020). These same years also saw a rise in high-profile cyberattacks. The Council of the EU later imposed sanctions on those behind the WannaCry and NotPetya attacks, showing how seriously these incidents were taken at EU level (Zygierewicz, 2020). At the same time, the average cost of cybercrime rose, and so did global cybersecurity spending, between 2017 and 2018 (Zygierewicz, 2020).


This sequence of events offers a plausible explanation for the cluster of new companies founded around 2016–2017 in the present sample. Regulatory pressure

and greater awareness of cyber-threats seem to have pushed firms to enter the market ahead of GDPR enforcement. The later decline in new firm creation after 2018 may reflect market saturation, as the first wave of compliance-driven entrants was absorbed. It may also reflect the reporting-lag effect noted earlier.

4.3.2. Closure trend

On the other hand, the closure trend in figure 4.2.3 displays a markedly different pattern from the founding distribution. Closures remained negligible throughout the 2010s, with no more than three firms closing in any given year up to 2017. From 2018 onwards, the number of closures began to rise steadily, reaching 14 in 2019 and 18 in both 2020 and 2021. This upward trend accelerated sharply thereafter, with closures climbing to 24 in 2023 and peaking at 25 in 2025, the highest annual figure in the sample. The sharp fall recorded for 2026 should again be read with caution, given the likelihood of reporting lag for the most recent period.

Figure 4.3: Closure trend of companies



This pattern of accelerating closures from 2022 onwards fits a broader macro-financial context. This context is well documented in the academic literature on venture capital.

Based on this trend, there are three studies that highlight the same mechanism. Ma and Zimmermann (2023) find that rate hikes are followed by a 25 % fall in venture capital investment over one to three years, an effect strongest among early-stage start-ups. They note a roughly 30 % annual decline across most sectors since 2022. Bellavitis, Fisch and Vismara (2023), using data from 31 countries (Europe, North America, Latin America, Asia-Pacific, Middle East and Africa), show that VC fundraising rises when interest rates turn negative, implying the reverse as rates climb.¹ Allen (2025) frames this within a longer cycle: loose monetary policy between 2009 and 2021 fed a VC bubble, which the 2022 tightening then unwound, with knock-on effects for banks serving venture-backed firms.

Together, these studies offer a plausible explanation for the closure pattern in the present sample. New firm creation peaked around 2017 and was reinforced by favourable financing conditions in 2020–2021, producing a cohort of firms especially exposed to the funding contraction that followed the 2022 rate rises. The rise in closures between 2022 and 2025 fits this timeline, as affected firms gradually ran out of cash once follow-on capital became harder to raise, consistent with the liability-of-newness and liability-of-adolescence patterns typically found in venture-backed technology sectors.

4.3.3 Entry cohorts

In light of the evidence presented above, the analysis proceeds on the basis of age cohorts. This makes it possible to identify the distinct periods in which firms in the dataset are concentrated. These cohorts are then used later as fixed effects. There are four classes: 2009–2012, 2013–2017, 2018–2022, 2023–2026. This periodisation is not arbitrary. It reflects the distinct regulatory and financial phases identified earlier:

- I. The *2009–2012* cohort corresponds to a period of modest and stable firm creation, prior to the main wave of EU cybersecurity regulation;

¹ Argentina, Australia, Austria, Belgium, Canada, Chile, Colombia, Czech Republic, Denmark, France, Germany, Greece, Hungary, Ireland, Israel, Italy, Japan, Mexico, Netherlands, New Zealand, Norway, Poland, Portugal, Russia, South Africa, South Korea, Spain, Sweden, Switzerland, United Kingdom, and the United States.

- II. The 2013–2017 cohort captures the sharp rise in foundation activity that led up to the 2017 peak. This period coincides with the build-up to the NIS Directive and the GDPR, as well as the WannaCry and NotPetya attacks, which together raised regulatory pressure and risk awareness across Europe;
- III. The 2018–2022 cohort spans the period of gradual decline in new firm creation following GDPR enforcement, alongside the exceptional venture capital surge of 2020–2021 and the first signs of the 2022 monetary tightening cycle. This cohort is therefore likely to include firms exposed to both the most favourable and the most adverse financing conditions in the sample;
- IV. The 2023–2026 cohort represents the most recent firms in the dataset. It reflects the post-tightening environment, in which venture capital contraction is most pronounced, while also being subject to a degree of reporting lag, as noted previously.

Grouping firms into these four cohorts therefore allows the subsequent econometric analysis to control for period-specific shocks (regulatory, financial, or macroeconomic) that would otherwise be difficult to disentangle from firm-level characteristics.

4.3.4 Lifespan and survival

The descriptive analysis on the full sample of 1,227 companies, of these firms, 159 (13.%) had closed or been acquired by the end of the analysis period, while 1,068 (87%) remained active.

Table 4.3: Summary statistics on lifespan and survival

	N	Mean (lifespan)	SD	Min	Max
Operating status					
Closed	159	6.434	3.721	1	17
Active	1068	8.822	4.082	2	17

A clear divergence emerges between the two groups. Closed and acquired firms show a mean lifespan of 6.43 years (SD = 3.72) and a median of 6.00 years. A quarter of these failures occur within the first four years of operation. This fits the "liability of newness" and "valley of death" concepts discussed earlier, whereby a venture's early years carry the greatest risk of exit.

Active firms tell a different story. Their mean duration is considerably longer, at 8.82 years (SD = 4.08), with a median of 9.00 years. Three-quarters of these firms have already operated for at least six years, the same point at which closure risk appears to fall away most sharply among the failed cohort.

Both distributions are mildly right-skewed (0.51 for closed firms, 0.28 for active firms) and platykurtic. This points to a modest cluster of long-lived outliers at the upper end, rather than a sharply asymmetric distribution overall.

This difference in mean duration between the two groups is *statistically significant*. A two-sample t-test confirms that active firms operate, on average, 2.39 years longer than closed or acquired firms ($t = -6.96$, $p < 0.001$, 95% CI: -3.06 to -1.71), providing formal statistical support for the descriptive pattern observed above.

Table 4.4: T-test: Firm Lifespan by Status

	obs1	obs2	Mean1	Mean2	dif	St Err	t value	p value
duration by K ² : 0 1	159	1068	6.434	8.822	- 2.388	.343	-6.95	0

To formalise this analysis within a survival-time framework, the data were declared using the `stset` command, with firm closure or acquisition ($K = 0$) defined as the failure event. This yielded the full sample of 1,227 observations corresponding to 159 recorded failures and a total of 10,445 firm-years at risk, with observation spanning from $t = 0$ to a maximum of $t = 17$ years. This confirms that the dataset provides a robust and complete base of person-time exposure for subsequent

² K=operating status

survival analysis, with a failure rate of approximately 15.2 closures per 1,000 firm-years.

Figure 4.4: *stset* duration, within 6 years threshold




Figure 4.4 indicates that survival in *the sample is not randomly distributed over time*: firms that cross the six-year mark exhibit a markedly higher and statistically significant probability of remaining active, supporting the treatment of age-related thresholds as a meaningful covariate in the subsequent regression analysis.

4.4 Firms and founders characteristics

This section is grounded in the entrepreneurship and venture capital literature on founding structure and founder background, both of which are established predictors of funding outcomes. Building on this literature, two exploratory firm- and founder-level variables are constructed and tested in the present dataset.

The first variable concerns founding structure. The literature on this point is mixed. Greenberg and Mollick (2018) studied thousands of Kickstarter-funded ventures. They found that solo-founded businesses tended to survive longer than team-founded ones. Solo founders also generated more revenue, particularly compared

with two-person teams. Evidence from traditional venture-backed samples points the other way, however. Solo founders raise substantially less capital than founding teams. They also close a smaller share of funding rounds, relative to their share of new incorporations.

This contradiction is usually explained on the basis of investor bias. Investors often treat the presence of a co-founder as a heuristic for market validation: that is, a mental shortcut used to infer the quality of a venture in the absence of verifiable performance data, rather than as direct evidence of superior future performance.

This heuristic reflects investors' inability to easily assess the true viability of an early-stage idea; as a result, they tend to rely on observable proxies, such as team composition, as signals of legitimacy. The presence of a co-founder is interpreted as a form of informal, peer-level endorsement of the business idea, rather than as a demonstrated outcome.

It is this reliance on an indirect signal, rather than on verified performance, that the literature identifies as a potential source of bias in capital allocation — one that may explain why VC-backed teams outperform solo founders in fundraising terms even where survival and revenue data suggest the opposite. Given this ambiguity, a dummy variable is included in the present analysis to distinguish solo founders from founding teams, making it possible to test which pattern, if either, holds within the cybersecurity sample.

The second variable concerns founders' educational and professional background. Unlike founding structure, the evidence here is considerably more consistent. Hsu (2007) studied early-stage technology start-ups and found that prior founding experience increased the likelihood of securing venture capital funding.

This effect was strongest for founders with financially successful prior experience, and it also increased the resulting valuation of the venture. Hsu explains this through the accumulation of human and social capital. Subsequent research supports this finding: founder education is consistently used by investors as a signal, particularly when a firm has no performance history of its own (Amit et al.,

1995; Ko and McKelvie, 2018). Higher levels of education are associated with a greater likelihood of securing funding. This is generally explained through human capital theory, whereby education and relevant professional experience act as observable proxies for founder quality, mattering most when no operating track record is yet available. On this basis, founders' educational and professional background is likewise encoded as a dummy variable.

Both variables are included in the logit regression as exploratory covariates, alongside the other explanatory factors considered in the model, in order to test whether founding structure and founder background are themselves associated with the probability of securing follow-on funding.

4.4.1 Solo Founders versus Founding Teams

Central to the regression is the dummy variable distinguishing solo founders from founding teams. This variable helps to establish whether team-led firms attract significantly more funding than those founded by a single individual, and whether this difference is statistically significant.

Figure 4.5: distribution of solo founder vs founding teams




Figure 4.5 presents the distribution of founding structure across the full sample of 1,227 startups. The two categories are close to evenly split, with solo founders

representing a marginal majority of the sample and founding teams accounting for the remainder.

This near-even split is itself a noteworthy descriptive finding in itself. It suggests that, within the cybersecurity sector, solo founding is not a marginal phenomenon but a genuinely common pathway into entrepreneurship, comparable in scale to team-based founding. This is broadly consistent with recent evidence from other technology sectors, where solo-founded ventures have been shown to account for a substantial and growing share of new incorporations, even as the venture capital industry continues to display a documented preference for founding teams at the point of investment.

This balanced distribution provides a solid empirical basis for the subsequent analysis. With both categories well represented in the sample, the logit regression is well positioned to detect any statistically significant difference in funding outcomes associated with founding structure.

4.4.2. Founders' educational background

This sub-section examines the founders' educational background, coded as a dummy variable distinguishing degrees coherent with the cybersecurity domain from those that are not. Before turning to the classification criteria, one methodological point requires clarification.

Where a firm was founded by a team rather than a single individual, the educational background attributed to the firm was that of the most representative founder. This determination was made through individual analysis of each founder's curriculum vitae, generally identifying either the first-listed founder or, where distinguishable, the individual holding the role of Chief Executive Officer (CEO) or an equivalent administrative lead. This choice reflects standard practice in the entrepreneurship literature, whereby the lead founder's human capital is treated as the most relevant signal available to investors, given that this individual typically holds primary decision-making authority and is most visible in the firm's external representation (Hsu, 2007; Ko and McKelvie, 2018).


The classification of “industry coherent” versus “industry non-coherent” educational background draws directly on the sectoral taxonomy developed in Chapter 3. Degrees are coded as “industry coherent” if they fall within fields directly underpinning the technical and analytical competencies required across the four macro-clusters identified in the dataset, namely: computer science, information technology, cybersecurity or information security, software engineering, and closely related STEM disciplines such as mathematics, physics, and electrical or telecommunications engineering. This reflects the product-centric and highly technical nature of the *Core Cybersecurity & Risk Management* and *Software & Deep Tech* clusters in particular, where competencies such as cryptography, network architecture, and applied machine learning constitute the core value proposition of the firm.

Conversely, degrees in business administration, economics, law, the humanities, and other non-technical fields are coded as “industry non-coherent”, even where the founder subsequently built a career in the technology sector.

Turning to the descriptive evidence, Figure 4.6 plots the distribution of educational background coherence, disaggregated by founding structure. The pattern is strikingly consistent across both groups.

- I. Among solo founders, 93.1% hold a degree coded as coherent with the cybersecurity domain, against only 6.9% classified as “non-coherent”.
- II. Among founding teams, the distribution is nearly identical: 93.6% of representative founders hold a “coherent” educational background, compared with 6.4% who do not.

Figure 4.6: Educational background of team founder vs solo founder



This similarity across the two groups is a notable finding in itself. It suggests that, in this sample, the choice to found a firm solo or as part of a team has little to do with the founder's education. Technical training is the norm either way. This is consistent with the human capital and signalling literature discussed above, which treats founder education as a key credential in technology-heavy sectors, where investors know far less than founders about the true quality of the venture (Hsu, 2007; Ko and McKelvie, 2018). Much of the sample is highly specialised and product-focused, especially in the *Core Cybersecurity & Risk Management* and *Software & Deep Tech* clusters identified in Chapter 3. Given this, a strong prevalence of technically coherent educational backgrounds is consistent with our expectations.

At the same time, the limited variation observed in this dummy variable (with 'non-coherent' backgrounds representing only around 6.9 and 6.4 % of the sample in both groups) raises a methodological consideration to be addressed at the estimation stage. Given this imbalance, the explanatory power of educational coherence within the logit regression is likely to be limited, and results relating to this variable should be interpreted with appropriate caution.

4.4.3 Professional background


Alongside educational coherence, a second exploratory variable captures the founders' prior professional experience. Following the same logic applied to the educational background variable, this measure is based on the last role held by the representative founder immediately before establishing the firm: the first-listed founder in cases of a solo venture, or the lead founder identified through individual analysis of founders' CVs in the case of a founding team. The variable is coded as a dummy, distinguishing founders whose immediately preceding role was already within the cybersecurity sector from those whose prior occupation lay outside it.

This measure complements the educational variable discussed above by capturing a different, and arguably more direct, dimension of sector-specific human capital. While educational coherence reflects formal training, professional background reflects hands-on exposure to the sector's technologies, client relationships, and competitive dynamics immediately prior to founding. The entrepreneurship literature identifies this type of pre-entry industry experience as a further signal of founder quality, on the grounds that direct sectoral exposure equips founders with tacit knowledge and professional networks that are difficult to acquire through formal education alone (Hsu, 2007).

Figure 4.7 presents the distribution of this variable, again disaggregated by founding structure. The pattern closely mirrors that observed for educational background, albeit with somewhat greater variation.

- I. Among solo founders, 71.6 % had a last role within the cybersecurity sector prior to founding, against 28.4% coming from outside the industry.
- II. Among founding teams, the distribution is again very similar: 72.3% of representative founders held a prior role within the sector, compared with 27.7% who did not.

Figure 4.7: Professional background of team vs solo founder



As with educational background, the near-identical distribution across solo founders and founding teams suggests that prior sector experience, like educational coherence, does not appear to differ systematically according to founding structure within this sample. However, the somewhat larger non-coherent share observed here (around 28%, compared with roughly 7% for educational background) indicates that professional background offers greater variation for the subsequent logit estimation, and may therefore prove a more informative exploratory variable in the regression stage.


4.4.4 Founding rounds

The final firm-level characteristic considered in this section concerns the type of funding obtained by each company in the sample, providing a first indication of the extent to which firms in the cybersecurity sector rely on external capital versus self-financing.

Figure 4.8 shows the distribution of the sample across four funding categories. Bootstrapped firms represent the single largest group, accounting for 691 companies, equivalent to 56.3% of the sample. This indicates that more than half

of the firms in the dataset have relied primarily on internal resources rather than external investment, a pattern that is not uncommon in the broader start-up population but is nonetheless notable given the highly technical and often capital-intensive nature of cybersecurity ventures.

Figure 4.8: founding rounds achieved by companies



Among externally financed firms, the distribution is heavily skewed towards the earliest stage of the funding lifecycle. Early-stage firms constitute the second-largest group, with 374 companies (30.5%), while growth-stage firms account for a considerably smaller share, at 118 companies (9.6%). Late-stage firms represent the smallest category by a wide margin, with only 44 companies (3.6%) having reached this level of funding.

This pattern of progressive attrition across funding stages: from bootstrapped, to early-stage, to growth-stage, to late-stage, is consistent with the well-documented "funding funnel" observed in venture capital markets, whereby only a small fraction of externally financed firms successfully raise capital beyond the earliest rounds. Combined with the closure and survival patterns discussed above, this distribution suggests that the cybersecurity firms in the sample face a similar dynamic to that observed in the broader technology sector. Initial funding is comparatively

accessible, particularly at the bootstrapped and early stages, but subsequent rounds become progressively harder to secure, with fewer than one in ten sampled firms reaching growth-stage financing and fewer than one in twenty-five reaching late-stage financing.

4.4.5 Mode of exit

Focusing on the 159 closed firms identified in Section 4.3.4, this final sub-section examines *how* these firms exited the market, distinguishing outright shutdown from acquisition, and considers whether the two exit routes are associated with different funding profiles.

Table 4.5 reports the distribution of exit modes among the closed firms. Acquisition is by far the more common outcome, accounting for 113 firms (72%), while only 44 firms (28%) shutdown without being acquired. This distribution carries an important interpretative implication for the survival analysis mentioned above. Treating all 159 failures as equivalent would conflate two economically distinct outcomes.

Table 4.5: mode of exit among closed firms

	Freq.	Percent	Cum.
Shutdown	44	28.03	28.03
Acquired	113	71.97	100.00
			0
Total	157	100.00	

An acquisition typically reflects a degree of market validation. The firm's technology, team, or client base was judged valuable enough to be absorbed by another player, whereas a shutdown more plausibly reflects an inability to sustain operations. The fact that fewer than three in ten exits in this sample are outright shutdowns suggests that closure risk in the cybersecurity sector, at least as captured by this dataset, is driven primarily by consolidation rather than by outright commercial failure.

Table 4.6 turns to the funding profile of the closed firms as a whole. Two-thirds of closed firms (66%) were bootstrapped, while the remainder had raised some form of external capital: 25.2% at early-stage, 5% at growth-stage, and 3.8% at late-stage. This distribution is broadly comparable to that observed for the full sample in Section 4.4.4, suggesting that exit, in itself, is not obviously concentrated among either externally financed or self-financed firms.

Table 4.6: founding type among closed firms

	Freq.	Percent	Cum.
bootstrapped	105	66.04	66.04
early-stage	40	25.16	91.19
growth stage	8	5.03	96.23
late stage	6	3.77	100.00
Total	159	100.00	

A more informative picture emerges from Table 4.7, which cross-tabulates exit mode against funding type. Among shutdown firms, the funding profile is markedly narrower: 68.2% were bootstrapped and 31.8% had raised early-stage funding, with no shutdown firms recorded at growth-stage or late-stage. Among acquired firms, by contrast, the funding profile is more dispersed across the full spectrum, including 6.96% at growth-stage and 5.22% at late-stage, funding levels entirely absent among shutdowns. This pattern suggests that firms which succeeded in attracting later-stage capital were, within this sample, exclusively acquired rather than shut down, consistent with the interpretation that reaching more advanced funding rounds is associated with a stronger, more durable market position, even where the firm ultimately exits the market through consolidation rather than continued independent operation.

Table 4.7: mode of exit by funding type

mode exit	bootstrapped	early-stage	growth stage	late stage	Total
Shutdonwn	30	14	0	0	44
	68.18	31.82	0.00	0.00	100.00
Aquired	75	26	8	6	115
	65.22	22.61	6.96	5.22	100.00
Total	105	40	8	6	159
	66.04	25.16	5.03	3.77	100.00

First row has *frequencies* and second row has *row percentages*

4.5 Other variables

Additional variables considered in the dataset include the firms' geographical location, its size, and its revenue. These are employed as control variables in the second stage of the analysis.

- I. *Geographical location* captures regional variation in venture capital availability and ecosystem maturity, which may influence survival independently of firm-specific characteristics.
- II. *Firm size* proxies for resource endowment and organisational slack, both of which are associated with resilience in early-stage ventures.
- III. *Revenue* serves as a direct indicator of commercial traction and market validation, distinguishing firms with demonstrated demand from those still pursuing product-market fit.

Together, these variables allow the second-stage models to test whether the primary relationships identified in the first stage hold once broader structural and financial factors are accounted for. Their inclusion strengthens the robustness of the empirical strategy by controlling for confounding effects that could otherwise bias the estimated coefficients.

4.5.1 Revenues

The distribution of firm-level revenue is heavily skewed towards the lower end of the scale. The majority of firms in the sample report revenue below \$10 million,

with 43.60% falling under \$1 million and a further 48.66% falling within the \$1M–\$10M band. Together, these two categories account for over 92% of the sample, confirming that early-stage and small-revenue firms dominate the cybersecurity startup population under study. Firms with revenue exceeding \$100 million represent a negligible share (1.14% combined), consistent with the expectation that few ventures in this sector reach substantial commercial scale within the observation window. This concentration at the lower end supports the use of revenue as an explanatory variable, as it allows the second-stage models to control for the fact that most firms in the sample remain in early or intermediate stages of commercial maturity.

Table 4.8: distribution of revenue categories

RECODE of revenue cat v2	Freq.	Percent	Cum.
< \$1M	535	43.60	43.60
\$1M - \$10M	597	48.66	92.26
\$10M - \$100M	81	6.60	98.86
\$100M - \$500M	11	0.90	99.76
\$500M+	3	0.24	100.00
Total	1227	100.00	

4.5.2 Geographical distribution

The geographical distribution of firms reveals a clear concentration outside the North and East European regions. The largest share of firms (59.33%) is classified under "Rest of Europe", followed by North Europe (31.62%) and East Europe (9.05%). This pattern suggests that the sample is not evenly distributed across the continent, with North European ecosystems, likely reflecting more mature venture capital markets, accounting for roughly a third of observations, while East European firms remain a minority. This uneven distribution reinforces the rationale for including geographical pattern as a confirmation variable in the second-stage analysis, as regional disparities in ecosystem maturity and VC availability may otherwise confound the relationships identified in the primary models.

Table 4.9: distribution by geographic location

	Freq.	Percent	Cum.
East Europe	111	9.05	9.05
North Europe	388	31.62	40.67
Rest of Europe	728	59.33	100.00
Total	1227	100.00	

4.5.3 Company' size

Table 4.8 presents the distribution of company size, measured by number of employees. The results confirm the pattern already observed for revenue. Most firms fall into the two smallest size categories. 1–10 employees' category accounts for 37.16% of the sample. 11–200 employees' category represents the largest single group, at 58.84%. Together, these two categories comprise 96.01% of the total sample. This indicates that the vast majority of cybersecurity startups under study operate with fewer than 200 employees. Larger firms are comparatively rare. 201–1000 employees' category accounts for only 2.85% of the sample. 1000+ employees' category accounts for a marginal 1.14%. This distribution is consistent with the revenue findings presented in Section 4.5.1.

It reinforces the interpretation that the sample is dominated by early-stage and small-scale ventures. This concentration justifies the inclusion of company size as a confirmation variable in the second-stage models. It allows the analysis to account for firm-level resource constraints that may otherwise confound the estimated relationships.

Table 4.10: distribution of size firms' categories

	Freq.	Percent	Cum.
1-10	456	37.16	37.16
11-200	722	58.84	96.01
201-1000	35	2.85	98.86
1001+	14	1.14	100.00
Total	1227	100.00	

4.6 Summarize and method applied

4.6.1 Overview and econometric strategy

This chapter has presented the construction of the dataset and its descriptive properties. The final sample comprises 1,227 cybersecurity-related firms, covering the period 2009–2024. A set of variables was built to capture the dimensions relevant to firm survival.

- I. Founding structure is measured through *team_dummy*, and founder human capital through *prof_bg* and *edu_bg*: professional background and educational background.
- II. Sector positioning is reflected in *cluster_idx*, and firm age in the cohort variable *coorte*.
- III. Geographical location is recorded through *region_id*, and firm size and revenue are captured through *size_cat* and *revenue_cat*.
- IV. Financing activity is measured through *eq_dummy*, a binary indicator of whether a firm raised equity funding above the sample's average threshold.
- V. Firm survival is recorded through *K* (operating status), and, for the firms that exited the sample, *mode_exit* distinguishes shutdown from acquisition.

These variables will be employed in two-stage estimation strategy developed in Chapter 5.

The empirical strategy is structured around three research questions, each addressed through a distinct stage of the analysis.

- I. **RQ1.** What variables influence the likelihood of receiving funding for cybersecurity startups? This question will be addressed through a **logit model estimation**.
- II. **RQ2.** Conditional on being funded, what are the determinants of survival of cybersecurity startups? This question will be addressed through the **Cox proportional hazards model**, preceded by the Kaplan-Meier analysis that establishes its univariate groundwork.

- III. **RQ3.** Conditional on being funded, in what way do firms exit the market — through acquisition or failure? This question will be addressed through a **competing-risks study**, which disaggregates exit into its two constituent routes.

Each stage of the analysis builds on the preceding one. The logit model generates a predicted financing propensity score, which then enters the Cox model as an explanatory variable. The Cox model, in turn, establishes the general survival relationship that the competing-risks study further disaggregates by exit type. The three questions are therefore not independent enquiries, but sequential layers of the same empirical strategy.

4.6.2 Method applied

This sequence provides the theoretical foundation for the empirical strategy adopted in this dissertation, which will be outlined in Chapter 5:

- **Stage 1 – The financing threshold (Logit).** The literature reviewed in Section 2.2 suggests that a venture's founder-related characteristics (professional background, educational background, founding team composition) and structural characteristics (sector segment, cohort) jointly determine whether a venture succeeds in crossing the threshold of attracting external equity. This binary outcome is therefore modelled through a **logistic regression**. Rather than being carried forward into the survival model in its raw binary form, however, the fitted values of this first-stage model are retained as a continuous predicted propensity to attract equity financing, purged of the discrete threshold and expressed as a function of the venture's exogenous founder- and sector-related characteristics alone.
- **Descriptive validation (log-rank / sts tests).** Prior to the specification of a fully multivariate survival model, the theoretical relevance of the key covariates identified in Sections 2.1–2.3, financing status, founder team composition, professional and educational background, and mode of exit, is assessed independently through univariate tests of survivor function equality (log-rank tests).

This descriptive stage serves two purposes consistent with the theoretical framework developed above: first, it provides preliminary, model-free evidence on whether each construct discussed in the literature exhibits a distinguishable survival profile in isolation; second, and more importantly, it establishes the empirical rationale for proceeding to a multivariate specification. Where univariate tests fail to detect significant differences, this is consistent with the literature's repeated emphasis (Sections 2.1 and 2.2) that survival is a multidimensional construct shaped by the joint, and potentially confounded, influence of several structural factors simultaneously, a pattern that a bivariate test, by construction, cannot disentangle.

- **Stage 2 – The survival and hazard structure (Cox).** The predicted financing propensity derived from Stage 1 is introduced into a **Cox proportional hazards model**, together with structural control variables spanning firm size, revenue category, sector segment, age firm, cohort, and region.

This specification directly operationalises the theoretical mechanisms discussed in Section 2.1 (the Valley of Death and the liability of newness, here reframed in terms of firm size and structural cost exposure) and Section 2.3 (the accelerant role of VC capital), whilst allowing the marginal effect of each covariate on the instantaneous hazard to be estimated net of the confounding structural factors that the descriptive stage alone cannot isolate.

- **Extension: Competing risks by mode of exit.** Building on the theoretical distinction between "cash-out" and "flame-out" trajectories developed in Section 2.3.1, the *unified hazard specification* above is further disaggregated into a **competing-risks framework**, in which the market-exit event is separated into two mutually exclusive, cause-specific outcomes: operational shutdown and acquisition (M&A).

This extension is theoretically motivated rather than a purely technical refinement: the literature reviewed in Section 2.3 suggests that VC-backed ventures are pushed towards *polarised*, not merely *accelerated*, outcomes, and a single undifferentiated failure event is, by construction, unable to distinguish a venture accelerated towards a favourable exit from one accelerated towards insolvency.

The cause-specific decomposition therefore provides the appropriate empirical counterpart to the theoretical polarisation mechanism discussed in Section 2.3.1. It should be noted that the IPO pathway, whilst theoretically central to the "cash-out" narrative developed in Section 2.3.1, is not modelled as a separate competing risk in this dissertation, owing to the absence of publicly listed firms within the sample. The empirical analysis therefore treats acquisition as the relevant observable proxy for a successful, VC-aligned liquidity event, whilst shutdown captures the opposing "flame-out" trajectory. It has to be noted that the theoretical foundation is deliberately sector-agnostic, drawing on general entrepreneurial and corporate finance literature applicable to high-technology ventures at large.

However, several of the mechanisms discussed, the severity of the Valley of Death, the composition of the founding team, and the pathway to external financing, are likely to manifest with intensity within specific high-technology domains. The following chapter examines how these general dynamics manifest within the cybersecurity sector specifically, providing the sector-specific theoretical and empirical grounding necessary to justify the variable operationalisation adopted in the remainder of this dissertation.

CHAPTER 5: EMPIRICAL ANALYSIS

5.1 Overview

The empirical strategy adopted in this dissertation proceeds in two stages.

The first stage models the probability of receiving equity financing as a function of firms' characteristics measured at, or close to time of firms' founding. A Logit model estimates this probability. The dependent variable, *eq_dummy*, captures whether a firm secured equity funding. Predictors include: founding team composition, professional and educational background, cluster membership, and time-founding cohort. This stage does not aim at estimating survival. It aims at isolating the component of equity financing that is predictable from the initial conditions, before any post-founding performance can influence the outcome.

The second stage employs the fitted propensity estimated from this first model, *eq_dummy_hat*, as an instrument for the estimation of a survival *model*. This choice addresses a specific concern: equity financing is unlikely to be exogenous to firm performance. Firms that perform well may find it easier to raise equity, and firms that raise equity may perform better as a result. Estimating the effect of equity financing on survival directly would conflate these two directions of causality. Using the predicted propensity score, rather than the observed financing status, helps isolate the component of financing attributable to founding-stage characteristics.

Before turning to the multivariate survival model, the analysis presents univariate descriptive evidence through the *Kaplan-Meier estimator*. This non-parametric approach traces the probability of survival over time without imposing assumptions on the functional form of the hazard. It offers a first, transparent picture of how survival patterns may differ across firm characteristics and serves as a useful preliminary analysis for the subsequent Cox estimation.

The Cox proportional hazards model then forms the core of the second stage. It estimates the hazard of firm exit as a function of the instrumented equity financing

variable, together with a set of other covariates: firm age, size category, revenue category, cluster membership, cohort, and region. The model does not assume a specific baseline hazard function, which makes it well suited to a dataset with heterogeneous firm trajectories. Standard errors are estimated using a robust variance estimator, to account for potential heteroskedasticity within the sample.

Finally, the empirical analysis extends beyond a single, undifferentiated notion of "exit". Based on the current dataset, firm exit can take one of two forms: acquisition or shutdown. A competing-risks framework accommodates this heterogeneity. It models the hazard of each specific mode of exit separately, allowing the analysis to distinguish equity financing's association with a successful exit through acquisition from its association with closure. This distinction is central to interpreting the results in the sections that follow.

All in all, each stage addresses a specific methodological concern, moving from the selection into equity financing, through descriptive survival patterns, to a fuller account of how financing relates to different pathways out of the market.

5.2 Logit estimation

The dependent variable of interest, *eq_dummy*, is binary: a firm either secures equity financing, or it does not. Standard linear regressions are not well suited in this setting. A linear probability model can generate fitted values outside the [0,1] interval, and it assumes a constant marginal effect of each regressor across the full range of the predicted probability. Neither property is appropriate here, where the object of interest is a probability by construction.

The Logit model addresses both limitations. It rests on a latent variable formulation. An unobserved continuous variable, y_i^* , is assumed to capture the underlying propensity of firm i to attract equity financing:

$$y_i^* = \beta_0 + \beta_1 \text{team}_{\text{dummy}_i} + \beta_2 \text{prof_bg}_i + \beta_3 \text{edu_bg}_i + \beta_4 \text{cluster_idx}_i + \beta_5 \text{cohort}_i + \varepsilon_i \quad (1)$$

The observed binary outcome, eq_dummy_i , takes the value 1 if this latent propensity crosses an unobserved threshold, conventionally normalised to zero:

$$eq_dummy_i = \begin{cases} 1 & \text{if } y_i^* > 0 \\ 0 & \text{if } y_i^* \leq 0 \end{cases} \quad (2)$$

The error term ε_i is assumed to follow a logistic distribution. This assumption yields a closed-form expression for the probability of financing, given by the logistic (sigmoid) function:

$$P(eq_dummy_i = 1 | X_i) = \frac{\exp(X_i \beta)}{1 + \exp(X_i \beta)} = \Lambda(X_i \beta) \quad (3)$$

where $X_i \beta$ denotes the linear index of covariates. This functional form bounds the predicted probability between 0 and 1 for any value of the index, and it allows the marginal effect of each regressor to vary with the level of the predicted probability, unlike the linear probability model.

Estimation proceeds by maximum likelihood. The coefficients β are chosen to maximise the likelihood of observing the actual pattern of financing outcomes in the sample:

$$L(\beta) = \prod_{i=1}^N \Lambda(X_i \beta)^{eq_dummy_i} [1 - \Lambda(X_i \beta)]^{1 - eq_dummy_i} \quad (4)$$

The resulting coefficients are interpreted as changes in the log-odds of financing associated with a one-unit change in the corresponding regressor, holding the other covariates constant. A positive coefficient indicates that the characteristic is associated with a higher probability of financing; the exact magnitude of this association is non-linear and depends on the baseline level of the predicted probability.

Within the two-stage strategy, the fitted values from this model, $\hat{P}(eq_dummy_i = 1 | X_i)$, are retained as the propensity score. This score is subsequently used, rather

than the raw financing dummy, as the explanatory variable of interest in the Cox proportional hazards model.

Table 5.1 above reports the coefficients, standard errors, and significance levels for the full specification.

Table 5.1: Logit model

	Co ef.	St.Err.	t- valu e	p-value	[95% Conf	Inter val]	Sig
team_dummy	.436	.136	3.21	.001	.17	.703	***
prof_bg	.097	.184	0.53	.597	-.264	.458	
edu_bg	.112	.269	0.42	.678	-.416	.64	
Core cybersecurity~e	-.01	.188	-0.05	.96	-.378	.359	
IT Services & Prof~u	-.566	.253	-2.24	.025	-1.061	-.07	**
Residual category	1.299	.185	7.01	0	.936	1.662	***
base category: Software &~h							
base category: 2009-12							
2013-17	.365	.223	1.63	.102	-.073	.803	
2018-22	.351	.223	1.57	.115	-.086	.788	
2023-26	.633	.3	2.11	.035	.045	1.221	**
Constant	-1.975	.35	-5.64	0	-2.662	-1.288	***
Mean dependent var		0.292	SD dependent var		0.455		
Pseudo r-squared		0.105	Number of obs		1227		
Chi-square		138.645	Prob > chi2		0.000		
Akaike crit. (AIC)		1345.709	Bayesian crit. (BIC)		1396.832		

*** $p < .01$, ** $p < .05$, * $p < .1$

All in all, the model performs well overall. The Wald chi-square statistic is significant (138.65, $p < 0.001$). The covariates jointly explain a meaningful share of the variation in equity financing.

The pseudo R^2 of 0.105 is modest. This is typical for cross-sectional logit models of firm-level financing. Founder and sector characteristics are expected to capture only part of the variation. Two factors dominate the results: team composition and sector positioning. Founder background and founding cohort play a smaller role, as discussed below.

Two results stand out.

First, team composition is the strongest predictor of financing. Firms founded by a team are significantly more likely to receive equity investment than solo-founded firms ($\beta = .436$, $p = .001$). Investors appear to treat a plural founding structure as a positive signal. This holds independently of the founders' professional or educational profile.

Second, sector positioning matters, but its effect is uneven. Firms in *IT Services & Professional* are significantly less likely to secure financing than the base category, software & deep tech ($\beta = -.566$, $p = .025$). Firms in the *Residual category* show the opposite pattern: a significantly higher likelihood of financing ($\beta = 1.299$, $p < .001$). This is the largest effect in the model.

Among the cohort variables, only the most recent period is significant. Firms founded in 2023–26 show a markedly higher probability of financing than those founded in 2009–12 ($\beta = .633$, $p = .035$). This is consistent with the recent acceleration of early-stage investment in cybersecurity.

Founder background does not appear to drive financing outcomes. Neither professional background (*prof_bg*, $p = .597$) nor educational background (*edu_bg*, $p = .678$) is statistically significant. Once team structure and sector are accounted for, these individual-level characteristics add no further explanatory power.

Overall, the results show that financing decisions are not driven by a founder's individual profile. They are driven by structural choices and sector context. A team-based founding structure raises the estimated probability of financing by a similar margin to moving from the least to the most favourable sector cluster.

In practical terms, how a firm is founded, and where it is positioned, appear to matter more to investors than who founded it. This distinction is explored in more depth in Chapter 6, where these patterns are linked to the survival outcomes examined in the second stage of the analysis.

5.3 Kaplan-Meier analysis

Before moving to the estimation of the Cox proportional hazards model, this section presents a preliminary descriptive analysis using the *Kaplan-Meier estimator*. The aim is not to establish causal relationships. It is to provide a first, non-parametric picture of survival patterns employing the key variables identified in the first-stage logit model.

The Kaplan-Meier estimator is a non-parametric method for estimating the survival function from time-to-event data. It makes no assumption about the shape of the underlying hazard. Instead, it builds the survival curve directly from the observed exit events. The estimated probability of survival is updated each time a firm leaves the sample. This makes the method well suited to the present dataset, which includes censored observations: firms still active at the end of the observation window, whose true survival time remains unknown.

This analysis serves two purposes:

- I. first, it provides a useful descriptive check on whether the variables that predict financing, notably team composition and sector cluster, are also associated with observable differences in survival;
- II. second, it provides a natural bridge from the Logit model to the Cox model, which extends this descriptive evidence into a formal multivariate

framework capable of isolating the effect of each covariate while controlling for the others.

The choice to begin with this step follows a basic principle of survival analysis. A variable is more likely to be significant in a demanding multivariate model, such as the Cox, if it shows some bivariate link to the outcome.

The Kaplan-Meier framework, together with the log-rank test, provides this preliminary check. The log-rank test compares the observed number of failures in each group against the number expected under the null hypothesis of equal survival functions. A significant result points to divergence between groups. A non-significant result suggests that, on that dimension alone, the groups do not differ.

The survival-time data were set using *survival_time* as the analysis-time variable and *fail == 1* as the failure event. The sample includes 1,227 firms and 159 failures, with a total of 10,429 firm-years at risk. The last observed exit occurs at year 17. Table 5.2 summarises these baseline settings.

Table 5.2: survival-time data analysis

Statistic	Value
Total observations	1,227
Failures (<i>fail</i> == 1)	159
Total analysis time at risk (firm-years)	10,429
Last observed exit (years)	17

Against this baseline, four founder- and financing-related variables were tested: *eq_dummy*, *team_dummy*, *prof_bg*, and *edu_bg*. The results are summarised in Table 5.3.

Table 5.3: log-rank test results

Variable	$\chi^2(1)$	p-value
<i>eq_dummy</i>	0.01	0.9146
<i>team_dummy</i>	1.52	0.2175
<i>prof_bg</i>	0.23	0.6299
<i>edu_bg</i>	0.24	0.6218

None of the four tests reject the null hypothesis of equal survival functions. The result for *eq_dummy* is particularly close to zero ($\chi^2 = 0.01$, $p = 0.91$). Observed and expected failures are almost identical across the two groups of firms: funded and not funded.

Figure 5.1: KM survival estimates by funding status (*eq_dummy*)





Figure 5.1 shows two nearly overlapping survival curves. Funded and non-funded firms follow the same downward path throughout the observation window. There is no visible point of departure between the two lines, at any stage of firm life. This flat, overlapping profile is the graphical counterpart of the near-zero χ^2 statistic in Table 5.2.

The result for *team_dummy* is somewhat larger ($\chi^2 = 1.52$, $p = 0.22$). Solo founders show a marginally higher-than-expected number of failures. This deviation remains well within the range attributable to chance. Figure 5.2 illustrates this modest divergence.

Figure 5.2: KM survival estimates by team coposition (*team_dummy*)



Moreover, Figure 5.2 shows a small separation between the two curves, emerging from around year eight onwards. Team-founded firms sit marginally above solo-founded firms for the remainder of the observation period. The gap is visible but narrow. It does not reach the width associated with a statistically robust divergence, consistent with the non-significant test result.

Additionally, *prof_bg* and *edu_bg* show near-identical survival curves and comparably weak test statistics. Figures 5.3 and 5.4 present these two variables in turn.

Figure 5.3: KM estimates by professional background (*prof_bg*)




Figure 5.3 shows two curves that move almost in lockstep across the full time range. Firms whose founders hold a relevant professional background show no discernible survival advantage over those without one. The lines cross more than once, a further sign of the absence of any consistent ordering between the two groups.

Figure 5.4: KM estimates by educational background (*edu_bg*)




Figure 5.4 shows a similar pattern for educational background. The two curves track closely for the first ten years. A small divergence appears towards the end of the observation window, where firms without a relevant educational background show marginally higher survival. Few firms remain at risk at this late stage. This late-

stage divergence should therefore be read with caution, rather than as evidence of a substantive effect.

Taken together, these results confirm the diagnostic conclusion set out above: *none of the founder- or financing-related characteristics show a detectable bivariate association with survival*. This should not be read as a null result in the ordinary sense. It is better understood as a diagnostic finding. Two explanations are consistent with this pattern:

- I. either these characteristics genuinely carry little independent explanatory power over survival,
- II. or their effects are masked at the bivariate level by confounding structural factors, such as sector or firm age, which the log-rank test cannot condition on.

This ambiguity is precisely what motivates the move to the Cox model, where the effects of these variables can be re-estimated jointly, net of the structural controls that a univariate test is unable to accommodate.

5.4 Survival model: COX proportional hazards model

Before moving to the estimation of the Cox model, the predicted propensity score, *eq_dummy_hat*, was checked for basic diagnostic soundness. Table 5.4 reports summary statistics for this variable.

Table 5.4: summary statistics predicted (*eq_dummy_hat*)

Variable	Obs.	Mean	Std. dev.	Min	Max
<i>eq_dummy_hat</i>	1,227	.2918	.1640	.0731	.6462

Two features of this distribution are worth noting.

- I. First, the mean predicted probability, 29.18%, matches the actual share of financed firms in the sample almost exactly. This is consistent with correct calibration of the first-stage logit model.

- II. Second, the predicted values remain well within the (0,1) interval, ranging from 7.3 to 64.6%. There is no evidence of degenerate values clustering near 0 or 1.

The lowest predicted propensity identifies an unfavourable profile: a startup with just one founder, in the *IT Services* cluster, belonging to the earliest cohort. The highest predicted propensity corresponds to the opposite: a startup with a multi-founder team, in the *Residual* cluster, from the most recent cohort. This spread suggests that the score behaves as a reasonably well-calibrated predictor, rather than an artefact of overfitting.

Before employing this predicted in the Cox estimation, a further methodological step was required. The continuous covariates, namely firm age and the predicted financing propensity *eq_dummy_hat*, were centred on their sample means. This step is standard practice when the aim is to derive baseline survival and cumulative hazard functions with a meaningful economic interpretation. Without centring, these baseline functions, generated through the *basesurv* and *basechaz* options, would describe a theoretical firm with all covariate values set to zero. Such a firm has no economic meaning: a firm of zero age, with zero financing propensity, does not exist in the sample.

Centring resolves this. It shifts the intercept of the proportional hazards model, so that the baseline curves instead trace the trajectory of the ‘average firm’ in the sample: a firm of approximately nine years of age, with a predicted financing propensity of 29.18%.

This adjustment does not alter the model's coefficients or their statistical significance. Its purpose is purely one of interpretability. It ensures that the baseline survival and hazard curves presented below can be read as describing a realistic, average firm in the dataset, rather than an abstraction with no counterpart in the data.

The Cox model extends beyond the covariates used in the first-stage logit. Firm age (*age_firm_c*) is introduced as a continuous control, alongside the centred financing

propensity discussed above. Age captures any residual association between firm maturity and exit risk that operates independently of financing status.

Four categorical variables are introduced as fixed effects: firm size (*size_cat*), revenue category (*revenue_cat*), sector cluster (*cluster_idx*), and founding cohort (*coorte*). Firm size and revenue category did not appear in the first-stage logit, since neither variable is determined at, or close to, founding. Both evolve over the life of the firm and are therefore unsuitable as exogenous predictors of the initial financing decision.

Within the Cox model, however, they serve a different purpose: they act as structural controls, absorbing heterogeneity in exit risk associated with firm scale and revenue performance, independently of the financing propensity score.

Sector cluster and founding cohort appear in both stages of the analysis, but their role differs between them. In the first stage, they help to explain the likelihood of financing. In the second stage, they are treated as independent structural determinants of survival, operating through channels distinct from financing, such as sector-level competitive intensity or the maturity of the investment environment at the time of founding. Their inclusion in both models does not therefore imply redundancy but reflects a change in their theoretical role across the two stages of the analysis.

A regional fixed effect (*region_id*) is also introduced to account for any systematic geographic variation in exit risk not captured by the remaining covariates.

5.4.1 Centred regression

Together, these controls allow the effect of the predicted financing propensity to be estimated net of firm-level structural heterogeneity, addressing the confounding concern raised at the Kaplan-Meier analysis.

Table 5.5: centred regression

_t	Haz. ratio	Robust St.Err.	z	P> z	[95% Conf.	Interval]	Sig
eq_dummy_hat_c	.224	.449	-0.75	.456	.004	11.406	
age_firm_c	1.019	.061	0.32	.751	.906	1.146	
size_cat (base category: 1-10)							
11-200	1.625	.305	2.59	.010	1.125	2.346	**
201-1000	4.309	1.824	3.45	.001	1.879	9.880	***
1001+	5.836	3.713	2.77	.006	1.677	20.310	***
revenue_cat (base category: < \$1M)							
\$1M - \$10M	.498	.086	-4.04	.000	.355	.698	***
\$10M - \$50M	.758	.234	-0.90	.368	.414	1.387	
\$50M - \$100M	.375	.444	-0.83	.407	.037	3.822	
\$100M - \$500M	.377	.325	-1.13	.257	.070	2.040	
\$500M - \$1B	1.07e-17	1.09e-17	-38.38	.000	1.45e-18	7.86e-17	***
\$1B - \$10B	134.962	97.813	6.77	.000	32.607	558.619	***
cluster_idx (base category: Software & Hardware)							
Core	1.543	.321	2.09	.037	1.027	2.320	**
cybersecurity & Risk							
IT Services & Professional	.831	.252	-0.61	.541	.458	1.506	
Residual	.910	.602	-0.14	.887	.249	3.329	
category coorte (base category: 2009- 12)							
2013-17	1.713	.617	1.50	.135	.846	3.468	
2018-22	2.891	1.735	1.77	.077	.892	9.376	*
2023-26	1.583	1.707	0.43	.670	.191	13.107	
region_id (base category: East Europe)							
North Europe	1.184	.359	0.56	.578	.653	2.145	
Rest of Europe	1.227	.350	0.72	.474	.701	2.147	
Time at risk		10.429	Number of obs				
Wald Chi-square		2823.38	Prob > chi2				
*** $p < .01$		** $p < .05$			* $p < .1$		

Table 5.5 reports the results of the Cox proportional hazards model, with the continuous covariates entered in *centred form*. The coefficient on the predicted financing propensity, *eq_dummy_hat_c*, is small in magnitude (HR = 0.224). It

does not reach conventional levels of statistical significance ($p = 0.456$), with a wide 95% confidence interval $[0.004, 11.406]$ that spans several orders of magnitude. On this evidence, financing propensity shows no reliable association with overall survival once the structural controls are accounted for.

Firm age (*age_firm_c*) is likewise non-significant ($HR = 1.019$, $p = 0.751$). Moreover, section 5.5 examines the baseline hazard function directly and considers whether its shape over time is consistent with this result.

Firm size shows the clearest and most robust pattern in the model. The hazard of exit rises monotonically with size, from $HR = 1.625$ ($p = 0.010$) for the 11–200 employee band, to $HR = 4.309$ ($p = 0.001$) for the 201–1000 band, up to $HR = 5.836$ ($p = 0.006$) for the largest band (1001+). Larger firms in this sample therefore face a markedly higher hazard of exit than the smallest firms, a pattern discussed further in relation to mode of exit below.

Revenue category presents a more heterogeneous picture. Firms earning between \$1M and \$10M show a significantly lower hazard than the smallest revenue band ($HR = 0.498$, $p < 0.001$), consistent with the idea that early revenue traction is protective against exit.

The intermediate revenue bands (\$10M–\$50M, \$50M–\$100M, \$100M–\$500M) are not statistically significant, though their point estimates are consistently below 1.

Two categories warrant particular caution:

- I. The \$500M–\$1B band returns a hazard ratio of essentially zero ($1.07e-17$), a numerically degenerate estimate that reflects perfect or near-perfect prediction within a very sparsely populated cell rather than a genuine effect.
- II. Conversely, the \$1B–\$10B band returns an extreme hazard ratio ($HR = 134.962$, $p < 0.001$); given the very small number of firms in this category, this result should likewise be treated as a data limitation rather than substantive evidence of risk.

Sector effects are limited. Only the Core cybersecurity & Risk cluster reaches significance, showing a higher hazard than the base category, Software & Hardware (HR = 1.543, $p = 0.037$). Neither IT Services & Professional (HR = 0.831, $p = 0.541$) nor the Residual category (HR = 0.910, $p = 0.887$) differs significantly from the base.

Founding cohort and region contribute little explanatory power. Only the 2018–22 cohort approaches significance, at the 10% level (HR = 2.891, $p = 0.077$), suggesting a somewhat elevated hazard for firms founded in this period relative to the earliest cohort (2009–12); the remaining cohort dummies and both regional dummies fall well short of conventional thresholds.

Taken together, these results indicate that structural characteristics, principally firm size, and to a lesser extent revenue category and sector, are the dominant drivers of exit risk in the unified model, while financing propensity, the variable of central theoretical interest, shows no significant association with overall survival at this stage of the analysis.

5.4.2 Robustness check: uncentred regression

As a robustness check, the regression was re-estimated without centring the continuous covariates. Table 5.6 reports the results of this uncentred specification, alongside the same set of structural controls discussed above: firm size, revenue category, sector cluster, founding cohort, and region.

Table 5.6: Uncentred regression

_t	Haz. ratio	St.Err.	z	P > z	[95% Conf	Interval]	Sig
eq_dummy_hat	.229	.461	-0.73	.463	.004	11.75	
age_firm	1.01	.06	0.17	.863	.899	1.136	
base category:	1						
1-10							
11-200	1.631	.305	2.61	.009	1.13	2.355	***
201-1000	4.371	1.873	3.44	.001	1.887	10.124	***
1001+	5.931	3.798	2.78	.005	1.691	20.807	***
base category:	1						
< \$1M							
\$1M - \$10M	.498	.086	-4.03	0	.354	.699	***
\$10M - \$50M	.757	.236	-0.89	.371	.411	1.394	
\$50M - \$100M	.379	.45	-0.82	.414	.037	3.886	
\$100M - \$500M	.375	.325	-1.13	.257	.069	2.045	
\$500M - \$1B	0						
\$1B - \$10B	137.645	99.874	6.79	0	33.2	570.669	***
Core	1.549	.323	2.09	.036	1.029	2.332	**
cybersecurity~e							
IT Services &	.828	.251	-0.62	.534	.456	1.501	
Prof~u							
Residual	.903	.599	-0.15	.878	.246	3.312	
category							
base category:	1						
Software &~h							
base category:	1						
2009-12							
2013-17	1.718	.627	1.48	.138	.84	3.512	
2018-22	2.827	1.709	1.72	.086	.865	9.245	*
2023-26	1.532	1.659	0.39	.693	.184	12.793	
base category:	1						
East Europe							
North Europe	1.183	.36	0.55	.581	.652	2.146	
Rest of Europe	1.224	.35	0.71	.48	.699	2.144	
Mean dependent var		8.513	SD dependent var			4.114	
Pseudo r-squared		0.029	Number of obs			1227	
Chi-square		437.628	Prob > chi2			0.000	
Akaike crit. (AIC)		2072.739	Bayesian crit. (BIC)			2164.761	

*** $p < .01$, ** $p < .05$, * $p < .1$

Centring a continuous covariate is a linear rescaling and should, in principle, alter neither the coefficients nor their statistical significance in a standard Cox regression. The uncentred results are consistent with this expectation: the coefficient on the predicted financing propensity, *eq_dummy_hat*, remains small in magnitude and statistically non-significant (HR = 0.229, $p = 0.463$), closely matching the centred estimate reported above (HR = 0.224, $p = 0.456$). The near-

identical point estimates and confidence intervals across the two specifications indicate that the non-significance of financing propensity is not an artefact of centring, but a stable feature of the model.

The structural controls largely reproduce the pattern already established. *Firm size* shows a clear, *monotonic association with exit risk*, rising from a hazard ratio of 1.631 for the second size category ($p = 0.009$) to 5.931 for the largest ($p = 0.005$).

Read alongside the competing-risks results in Section 5.6, this gradient is more consistent with larger firms exiting predominantly through acquisition rather than through outright failure, than with any straightforward "liability of smallness" account.

Firm age (*age_firm*) remains non-significant ($HR = 1.01$, $p = 0.863$) in the uncentred specification as well, confirming the result obtained above.

Revenue category is more heterogeneous, but one result stands out on substantive grounds: firms earning between \$1M and \$10M face a markedly lower hazard than the smallest revenue band ($HR = 0.498$, $p < 0.001$), consistent with the idea that early revenue traction meaningfully lowers exit risk. The extreme hazard ratio recorded for the \$1B–\$10B band (137.6, $p < 0.001$) should not be read as a substantive finding; it almost certainly reflects the very small number of firms populating this cell and is better treated as a data limitation than as evidence of a genuine effect.

Sector and cohort effects remain modest. Only the *Core cybersecurity* cluster reaches significance ($HR = 1.549$, $p = 0.036$), a more muted result than the volatility often attributed to this sector in the literature. With the exception of the 2018–22 cohort dummy, which is marginally significant at the 10% level ($HR = 2.827$, $p = 0.086$), none of the cohort or regional dummies reach conventional significance thresholds.

All in all, this robustness check corroborates the structural findings of the centred model in full, including the stability of the non-significant financing-propensity estimate across specifications.

5.5 Baseline functions

The Cox model is semi-parametric. It leaves the baseline hazard function unspecified and estimates it non-parametrically, using the Breslow estimator. The effect of covariates is captured separately, through the proportional hazards structure. This means the baseline functions reported below do not assume any particular distributional form, such as Weibull or exponential. Their shape is recovered directly from the pattern of exits observed in the data.

Under the proportional hazard's assumption, the baseline also acts as a reference profile. The hazard for any other firm is obtained by scaling the baseline hazard by a constant factor. This factor depends on the firm's covariate values and is assumed constant over time.

Figure 5.5: baseline survival function, estimated for the average firm




Figure 5.5 shows a gentle and gradual decline. Survival probability remains above 95% for the first eight years, before easing further to approximately 89% by year seventeen. There is no evidence of a sharp drop at any point in the observation window.

Figure 5.6: baseline cumulative hazard function, estimated for the average firm




Figure 5.6 presents the mirror image of this pattern. The cumulative hazard rises at a broadly steady, near-linear pace, without discontinuities or sudden jumps, reaching a level of 0.11 by the end of the period.

Read together, the two baseline functions point to a degree of structural resilience in the average firm. Once extreme scaling trajectories and the highest-risk sector segments are set aside, the typical cybersecurity startup in this sample appears comparatively stable in its capacity to remain in the market.

These two results come from different parts of the Cox model, and it is worth being clear about the distinction.

The coefficient on *age_firm* answers one specific question: at any given moment, does an older firm face a different hazard of exit than a younger firm? The model holds time fixed and compares firms of different ages to one another. A non-significant coefficient means firms do not differ from each other because of their age.

The baseline hazard function answers a different question. It describes how the hazard of exit for a *single, typical firm* changes as that *firm_ages*, from founding onward. It holds the firm fixed and lets time move forward.

These two things could, in principle, point in different directions. A model could show no difference between older and younger firms at any given moment, while still showing that the hazard facing any one firm rises or falls sharply over its own lifetime. The two results only tell the same story when they agree with each other, not by definition.

In this case, they do agree. The age coefficient is close to 1 and non-significant, meaning age carries little weight in explaining differences between firms. The baseline hazard, in turn, rises gradually and without sharp turning points, meaning there is no strong age-related pattern to explain in the first place. Read together, both results are consistent with the same underlying picture: age plays a limited role in this sample, whether compared across firms or traced within a single firm over time.

Firms in this sample appear to face a fairly stable background risk once the initial founding period has passed. Risk accumulates gradually and near-linearly across the observation window, with no sharp early peak and no rising-then-falling pattern.

5.6 Mode of exit

Building on the findings of the unified Cox model, the analysis is now extended to a competing-risks framework. The aim is to examine the determinants of mode of exit in more depth.

Two things need to be noted before undertaking this analysis: . the competing-risks framework, and the method used to estimate it. The framework is conceptual. It states that shutdown and acquisition are two mutually exclusive outcomes. Once a firm experience one outcome, it can no longer experience the other.

This framework can be estimated in more than one way. One option is a multinomial logit. This option exit as a single choice between categories, made at one point in time. It does not use information on how long each firm survived before exiting.

The method used in this dissertation is different. It is the *cause-specific hazards approach*. This method keeps the time dimension. It asks not only which exit route a firm takes, but also when.

Under this approach, two separate Cox models are estimated, one for each exit route.

- I. In the first model, shutdown is the event of interest. Firms that exit through acquisition are censored at the point of acquisition.
- II. In the second model, the roles are reversed: acquisition is the event of interest, and firms that shutdown are censored at that point.

This approach is not chosen for convenience. *It is a standard way of estimating a competing-risks system when the outcome of interest is time-to-event.* The overall likelihood of a competing-risks system can be split into separate parts, one for each cause. Each part can then be estimated on its own, using a standard Cox regression, without losing the competing-risks structure.

This choice was made because the research question concerns both timing and outcome. The dissertation is interested in whether financing propensity, size, and sector affect *how quickly* a firm shuts down or is acquired, not only *whether* it does. A multinomial logit cannot answer this question, because it does not model time. The cause-specific Cox approach can, and it also stays consistent with the unified Cox model used earlier in the chapter.

Estimating two separate cause-specific models allows each covariate to have a different coefficient for each exit route. This is the key advantage over the unified model, where only one, averaged coefficient was estimated per covariate.

The purpose of this exercise is not purely descriptive. It is designed to revisit the results reported above in a more disaggregated form. Where the unified survival model produced an effect that seemed ambiguous, or difficult to interpret in economic terms, a cause-specific breakdown may show that this ambiguity was really two opposing effects cancelling each other out, rather than one weak effect.

The financing-propensity variable (*eq_dummy_hat_c*), together with size and sector, are of particular interest here. Size and sector both showed a significant, distinct relationship with overall survival in the unified model (Section 5.4). Financing propensity, by contrast, showed no significant association at all. Whether this null result reflects a genuinely weak relationship, or two offsetting cause-specific effects that cancel out in the unified model, is the empirical question this section takes up.

Figure 5.7: distribution of mode of exit




Figure 5.7 reports out the distribution of exit types across the sample. The large majority of firms, 87.2%, remain active at the end of the observation window. Among the firms that have left the market, acquisition is markedly more common than shutdown: 9.2% of the sample exited through acquisition, against 3.6% through shutdown.

This asymmetry, more than two acquisitions for every shutdown, provides an early indication that exit in this sector does not primarily signal operational failure. The cause-specific models below examine whether this pattern holds once firm-level covariates are considered.

This asymmetry between shutdown and acquisition is not evenly spread across the sample. Figures 5.8 and 5.9 break down mode of exit by size and revenue category and reveal a sharp underlying pattern.

Figure 5.8 shows a clear pattern. Shutdown is present, though modest, in the three smallest size categories, ranging from 4.6% to 8.6% of firms. In the largest size category (1001+ employees), however, shutdown is entirely absent.

Acquisition frequency rises with size overall, from around 4-5% in the smallest category (1-10 employees) to a peak of over 23% in the 201-1,000 employee band, before easing slightly to around 21% in the largest category (1001+).

Taken together, these two trends point in the same direction: *as firms grow larger, exit through acquisition becomes more common, while exit through shutdown becomes rarer, and eventually disappears altogether.*

Figure 5.8: mode of exit by size category³





Figure 5.9 shows an even sharper version of the same pattern. Shutdown occurs only among firms earning less than \$10M in revenue. From the \$10M–\$50M category upwards, no firm in the sample has shut down, regardless of how high revenue climbs. Acquisition, meanwhile, is present across every revenue category, including the very top of the distribution: the single firm in the \$1B–\$10B band was acquired. The \$500M–\$1B category is the one exception with no recorded exit of either kind, though this reflects the very small number of firms in that category, only two, rather than a substantive pattern.

³ Size 1: 1-10; size 2: 11-200; size 3: 201-1000; size 4: 1000+

Figure 5.9: mode of exit by revenue category



Read together, Figures 5.8 and 5.9 support a straightforward conclusion. Within this sample, outright failure is a phenomenon confined almost entirely to small, low-revenue firms. Once a firm grows past a certain scale, roughly \$10M in revenue or 200 employees, shutdown becomes vanishingly rare, and the only forms of exit still observed are acquisition or continued operation.

This pattern has a direct consequence for the cause-specific regression reported in Table 5.7, below. Several revenue and size categories contain zero shutdown events. With no shutdowns to explain in these cells, the Cox partial likelihood cannot provide precise coefficient estimate.

Table 5.7 reports these coefficients as they were estimated, but flags them explicitly, since Figures 5.8 and Figure 5.9 already explain why they arise and confirm that this is a real feature of the data, not an error in the model.

Table 5.7: cause-specific Cox model: mode of exit(acquisition- shutdown)

VARIABLE	SHUTDOWN MODEL	ACQUISITION MODEL
eq_dummy_h~c	-0.558 0.877	-2.181 0.370
age_firm_c	0.143 0.185	-0.050 0.479
size_cat: 1-10		
11-200	-0.156 0.614	0.910 0.000
201-1000	2.230 0.000	1.638 0.002
1001+	-38.521 0.000	2.503 0.000
revenue_cat		
\$1M - \$10M	-1.116 0.001	-0.490 0.023
\$10M - \$50M	-40.011 0.000	0.171 0.593
\$50M - \$100M	-40.587 0.000	-0.527 0.670
\$100M - \$500M	-40.856 0.000	-0.474 0.592
\$500M - \$1B	-2.051 0.050	-36.174 0.000
\$1B - \$10B	-39.327	6.400 0.000
cluster_idx		
Core cybersecurity	0.003 0.994	0.579 0.015
IT Services	-0.694 0.240	-0.125 0.731
Residual category	-0.142 0.905	-0.286 0.730
coorte		
2013-17	1.896 0.015	0.194 0.646
2018-22	2.942 0.015	0.388 0.577
2023-26	2.724 0.172	-0.458 0.725
region_id		
North Europe	0.599 0.340	-0.039 0.911

Rest of Europe	0.216	0.149
	0.730	0.638

Table 5.7 reports the *cause-specific Cox estimates for the two exit routes*. As anticipated above, several coefficients in the size and revenues variables are implausibly large in magnitudes (in the region of -38 to -41) paired with p-values of 0.000. These arise in exactly the categories identified in Figures 5.8 and 5.9 as containing zero shutdown events and should not be read as extreme effects. They are the expected consequence of estimating a hazard for an outcome that never occurs within that category.

The financing-propensity variable, *eq_dummy_hat_c*, is not significant in either cause-specific model (shutdown: $b = -0.558$, $p = 0.877$; acquisition: $b = -2.181$, $p = 0.370$). This result resolves the empirical question posed at the outset of this section. The null association observed for financing propensity in the unified model (Section 5.4) is not the product of two offsetting, cause-specific effects that cancel out when combined. It reflects a genuine and consistent absence of association with either exit route, whether shutdown or acquisition is considered on its own.

Firm age (*age_firm_c*) remains non-significant in both models, consistent with the pattern already established in Section 5.4.

Among the non-degenerate size and revenue coefficients that are not large in magnitude, size shows a positive and significant association with acquisition throughout ($p < 0.01$ in every band), reinforcing the pattern shown in Figure 5.8: larger firms are systematically more likely to exit through acquisition.

Revenue shows a comparable pattern for shutdown, where interpretable: firms in the \$1M–\$10M category experience a significantly lower hazard of shutdown than those in the smallest category ($b = -1.116$, $p = 0.001$), the only revenue coefficient in the shutdown model not affected by the zero-event problem described above.

Sector effects are limited to acquisition mode of exit. The Core cybersecurity cluster shows a significant positive association with acquisition ($b = 0.579$, $p = 0.015$) but

not with shutdown ($b = 0.003$, $p = 0.994$). Neither IT Services nor the Residual category have significant coefficients in either model. Founding cohort shows some association with shutdown, with both the 2013–17 and 2018–22 cohorts significantly more likely to shut down than the earliest cohort, though these estimates should be read alongside the small number of shutdown events overall. Region shows no significant effect in either model.

All in all, the cause-specific breakdown analysis, rather than overturns, the picture built up over the course of this chapter. *Financing propensity shows no significant relationship with survival at any level of analysis*, whether the outcome is defined broadly, as in the unified model, or disaggregated by exit route, as in this section. On this evidence, the predicted propensity to raise equity financing does not appear to be a meaningful determinant of how, or how quickly, cybersecurity startups in this sample leave the market.

On the other side, size and revenue are strongly associated with survival in the unified model, and the cause-specific results show why. Larger, better-capitalised firms are not simply less likely to exit in general, they are specifically less likely to shut down, and correspondingly more likely to exit through acquisition.

What appeared in Section 5.4 as a single, undifferentiated hazard of exit rising with firm size is, in fact, a result almost entirely driven by the hazard of acquisition. The near-complete absence of shutdown among larger and higher-revenue firms, as shown in Figures 5.8 and Figure 5.9, is the clearest single finding to emerge from this chapter.

Sector plays a more limited role throughout, and where it does matter, it does so specifically through acquisition rather than shutdown, a distinction that the unified model alone could not have revealed.

5.7 Overview of findings and discussion

The empirical results in this chapter can now be interpreted in relation to the three research questions guiding this thesis:

- I. What variables influence the likelihood of receiving funding for cybersecurity startups?
- II. Conditional on being funded, what are the determinants of survival of cybersecurity startups?
- III. Conditional on being funded, in what way firms exit the market, (whether through, acquisition, or failure)?

The determinants of equity financing have been estimated through the Stage 1 logit model; The determinants of survival conditional on financing have been estimated through the Stage 2 Cox model; How firms exit the market have been estimated with the competing-risks extension that disaggregates acquisition and shutdown as distinct outcomes.

5.7.1 Research question 1: determinants of equity financing

The Stage 1 logit model estimates the probability that a cybersecurity startup secures equity financing (*eq_dummy*). The explanatory variables are team composition, founder background, sector cluster, and founding cohort. The model performs adequately for an exploratory specification of this kind (*pseudo R*² = .105; $\chi^2(11) = 138.65, p < .001; n = 1,227$).

Team composition is the strongest predictor of financing probability. *Team_dummy* is positive and highly significant (coeff. = .436, *p* = .001). This finding is consistent with a signalling account of early-stage investment. Substantial information asymmetry exists between founders and investors at this stage. A financial track record is typically unavailable. Product-market validation is typically unavailable too. A structured founding team therefore functions as an observable, low-cost signal of organisational capacity, in line with the founding-team composition literature. Investors appear to price this structural signal more heavily than any individual founder's credentials.

This reading is reinforced by the individual-level variables. Founder professional background (*prof_bg*: coeff. = .097, *p* = .597) and educational background (*edu_bg*: coeff. = .112, *p* = .678) are both statistically indistinguishable from zero. Read

alongside the strength of *team_dummy*, this suggests that investment committees discount individual human capital once team structure is considered. A single well-credentialled founder does not appear to substitute for a functioning team. This sits in some tension with human-capital accounts of financing access, which would predict a more direct effect of founder credentials, independent of team structure.

Sector cluster shows a more selective pattern.

- I. Relative to the *Software & Deep Tech base category*, firms in IT Services & Professional Consulting are significantly less likely to obtain equity financing (coeff. = -.566, $p = .025$). This is consistent with the lower capital intensity of service-based business models. Such models are typically funded through revenue rather than external equity. They also offer a less attractive path to the scale-driven returns that investors seek.
- II. Firms in the *Core cybersecurity cluster* show no significant difference from the base category (coeff. = -.01, $p = .960$). On this evidence, a core cybersecurity focus is neither a financing advantage nor a disadvantage relative to broader deep-tech positioning. This is a more selective pattern than the credit-rationing literature discussed in Chapter 3 would suggest, which anticipates a more uniform financing disadvantage for less differentiated or less defensible sub-segments. Sub-sector specialisation within cybersecurity does not appear to function as a first-order screening criterion in this sample.
- III. The *Residual category* shows a markedly higher coefficient, and the largest effect in the model (coeff. = 1.299, $p < .001$). Given the heterogeneous composition of this category, this result is treated with caution. No substantive economic interpretation is offered here (see Section 6.3).

Cohort effects are largely flat across most of the sample period. One exception stands out.

- I. Firms founded in 2023–26 are significantly more likely to secure equity financing than the 2009–12 base cohort (coeff. = .633, $p = .035$).

II. The intermediate cohorts, 2013–17 and 2018–22, show no significant difference.

This is consistent with the hot/cold financing-cycle literature and points to more favourable, or more active, financing conditions in the most recent period, plausibly linked to heightened investor attention to cybersecurity following a period of intensified cyber-risk exposure.

The cross-sectional design cannot separate a genuine cohort effect from a broader market-cycle effect on financing supply, however.

Taken together, the results indicate that financing decisions are shaped primarily by structural factors, not by individual founder characteristics: a team-based founding structure raises the estimated probability of financing by a margin comparable to moving from the least to the most favourable sector cluster.

Who founded the firm appears to matter less to investors than how it was founded, and where. This is broadly consistent with the founding-team literature, which treats team composition, rather than individual founder pedigree, as the central signal in early-stage investor screening.

5.7.2 Research question 2: determinants of survival conditional on financing

Conditional on financing, survival is examined descriptively through Kaplan-Meier estimates before turning to the Cox regression. The descriptive evidence points to firm size as the dominant source of differentiation in survival outcomes, with the baseline hazard function offering a further, more general insight into the shape of exit risk over time.

Figure 5.10: KM conditional survival by firm size




Figure 5.10 reports Kaplan-Meier survival curves by size category. The four curves track closely together in the early years of observation. They begin to separate from around year six onwards.

- I. Size category 4 (1001+) breaks away first, falling to approximately 69% survival by year eight.
- II. Size category 3 (201-1000) follows a broadly similar path, falling to around 61% by year ten and below 55% by year thirteen.
- III. Size categories 1 (1-10) and size 2 (11-200) remain close together throughout the observation window, ending at approximately 80% and 77% survival respectively.

This descriptive pattern is consistent with the hazard ratios reported in the Cox model: *the overall exit rate rises as firm size increases, and the rise is steepest for the largest firms in the sample. At this stage, however, "exit" still pools two very different outcomes together, acquisition and shutdown. Section 5.6 and Section 5.7.3 unpack this distinction directly, and the result changes the interpretation substantially.*

This pattern runs counter to the classical liability-of-smallness argument in industrial economics. That argument treats micro-firms as the most exit-prone segment of any population, owing to their limited resources and thin capital buffers.

The reverse holds here, at least on the surface: larger firms exit at a higher rate, not a lower one. Taken at face value, this could be read as scale itself introducing vulnerability, rather than protecting against it.

The cause-specific results reported in Section 5.6 show that this reading does not hold. Once shutdown and acquisition are separated, the higher exit rate among larger firms is driven almost entirely by acquisition, not by failure.

Larger cybersecurity firms in this sample are not more likely to shut down; they are more likely to be acquired. The Kaplan-Meier pattern shown here should therefore be read as an early signal of this acquisition-driven dynamic, not as evidence that scale itself is a source of operational fragility.

5.7.3 Research question 3: firms' exit route (conditional on being funded)

Research Question 3 asks, conditional on being funded, in what way cybersecurity startups exit the market: through shutdown or through acquisition. Section 5.6 has already shown that these two outcomes are asymmetric. This section uses the cause-specific Cox estimates in Table 5.7 to identify which firm characteristics are associated with each exit route.

Financing propensity plays no role in determining the mode of exit. The predicted propensity to raise equity financing (*eq_dummy_hat_c*) is not significant in either the shutdown model ($b = -0.558$, $p = .877$) or the acquisition model ($b = -2.181$, $p = .370$). This mirrors its null association with overall survival in the unified model (Section 5.4) and rules out the possibility that this null result was concealing two offsetting effects. Financing propensity, as measured here, is simply not a driver of exit in this sample, whichever form that exit takes.

Size and revenue, by contrast, are the clearest determinants of the exit route, and they point in opposite directions for the two outcomes. On the acquisition side, the


association with size is positive, significant, and strengthens monotonically: 11–200 employees ($b = 0.910$, $p < .001$), 201–1,000 employees ($b = 1.638$, $p = .002$), and 1,001+ employees ($b = 2.503$, $p < .001$). Larger firms are systematically, and increasingly, more likely to exit through acquisition, consistent with the pattern already shown in Figure 5.8, where acquisition rates rise from around 4% in the smallest size band to over 21% in the largest.

On the shutdown side, risk is concentrated at a different point in the size distribution. The Kaplan-Meier estimates of bankruptcy-free survival by firm size, showed in Figure 5.11, represent medium-sized firms recording by far the steepest decline, falling to approximately 85% by year four and to around 57% by year eight, consistent with the significant, positive shutdown coefficient for the 201–1,000 employee band in Table 5.7 ($b = 2.230$, $p < .001$).

Micro and small firms show a much shallower decline over the same period, remaining above 94% survival throughout. Large firms (1,001+ employees) record no bankruptcy events at all across the observation window: their survival curve remains flat at 1.00, which is the source of the degenerate coefficient reported for that band in the shutdown model ($b = -38.521$).

As Section 5.6 makes clear, this coefficient reflects the complete absence of shutdown events in that category rather than a genuine protective effect and should not be read as evidence that the largest firms are immune to failure; it simply means none failed within this sample's observation window.

Figure 5.11: bankruptcy risk



Revenue reinforces the same picture where the data allow interpretation. Firms earning \$1M–\$10M face a significantly lower shutdown hazard than the smallest revenue band ($b = -1.116$, $p = .001$), and no firm earning above \$10M in this sample has shut down at all, consistent with Figure 5.9. Beyond this single interpretable coefficient, the remaining revenue bands in the shutdown model are affected by the same zero-event problem as the largest size band and are not treated as genuine estimates.

Sector effects are confined to the acquisition route. Firms in the Core cybersecurity cluster face a significantly higher hazard of acquisition than the base category ($b = 0.579$, $p = .015$, corresponding to an approximately 78% higher hazard), but show no distinguishable shutdown risk ($b = 0.003$, $p = .994$). This is consistent with an active market for acquiring frontier cybersecurity capability, rather than with any greater operational fragility among Core cybersecurity firms. Neither IT Services & Professional Consulting nor the Residual category reaches significance in either cause-specific model, so no directional claim is warranted for either.

Taken together, these results answer the third research question directly. Conditional on being funded, mode of exit is not driven by financing propensity at all; it is driven overwhelmingly by scale.

Firms that grow beyond a modest size or revenue threshold do not exit less often. They simply exit differently. Their exits become almost exclusively acquisitions, not shutdowns. Outright failure stays concentrated among smaller, lower-revenue firms. It peaks specifically among medium-sized firms.

Sector adds a further, narrower layer to this picture. Core cybersecurity firms face a higher hazard of acquisition specifically, not of failure.

Read alongside the findings for Research Questions 1 and 2, this points to a clear conclusion. The same characteristics that attract financiers to a cybersecurity startup also make it an attractive acquisition target later on. For most firms that exit in this sector, exit is a market outcome. It is not a business failure.

5.8 Final discussion: confronting the evidence with the literature

Chapter 2 identified three theoretical strands. They run sequentially, from a venture's initial conditions, through founder-related financing decisions, to its ultimate exit trajectory. This section returns to each strand. It confronts the theory with the empirical results obtained in this thesis.

5.8.1 Organisational survival and the Valley of Death

Chapter 2 argued that survival is not a static condition. It is time dependent. It is shaped by a firm's stage of development. The earliest phase of a venture's life, the Valley of Death, was identified as the period of highest hazard. Resource scarcity drives this hazard. So does the rapid obsolescence of founding endowments. Founding-team composition was identified as the principal mitigating factor.

The empirical results partially support this argument:

- I. The founding-team result holds.
- II. Team composition is the strongest predictor of financing probability ($\beta = .436, p = .001$).

This is consistent with the theoretical claim that co-founding teams present a more credible proposition to investors than solo founders.

The age-dependent hazard is not confirmed, however. The literature predicts a high hazard immediately after founding. It then predicts a decline, as the firm builds absorptive capacity. The baseline hazard estimated in this thesis shows no such pattern. Risk rises gradually and near-linearly across the observation window. There is no early peak. There is no subsequent decline. This is a genuine departure from the literature.

The result carries a clear implication. For cybersecurity startups in this sample, the Valley of Death is not a discrete phase that firms simply cross or fail to cross. Vulnerability instead appears structural and persistent. It is rooted in size and sector, not in age.

This fits well with the sector-specific literature discussed in Chapter 3. The "cockroach" survival archetype (Ramsinghani, 2016, as cited in Zhang, 2016) emphasises capital efficiency and resilience over rapid scaling. A firm built for endurance, rather than for a rapid cash-out, would plausibly show exactly this kind of flat, persistent risk profile.

The size results add a further complication. The literature treats small, resource-constrained firms as the most vulnerable, owing to credit rationing and limited capacity to absorb fixed costs. The Cox results show the opposite ordering. Hazard rises monotonically with size, from 62.5% higher for firms with 11–200 employees ($HR = 1.625$, $p = .010$) to almost six times higher for firms above 1,000 employees ($HR = 5.836$, $p = .006$). This is not a simple rejection of the liability-of-smallness argument, though. Once exit type is disaggregated, the size effect is concentrated in the acquisition model. It is not present in the shutdown model. Larger cybersecurity firms are not failing at a higher rate. They are being acquired at a higher rate. The apparent reversal must be better read as a measurement artefact of the unified failure event, not as genuine evidence against the resource-based logic of the Valley of Death.

5.8.2 Founder background as a signal to investors

Chapter 2 argued that founder background is not an isolated determinant of survival. It functions as a strategic signal. It shapes a venture's capacity to attract capital, moderated by the founder's opportunity cost and by the macroeconomic cycle.

The results broadly support this signalling logic. They point to a different mechanism, however. Founder *professional background* ($p = .597$) and *educational background* ($p = .678$) are both statistically **indistinguishable** from zero, once team composition is accounted for. This does not confirm the human-capital account in the literature, which anticipates a direct effect of individual credentials. **It instead suggests that the relevant signal operates at team level**, not individual level. Investors appear to price organisational structure more heavily than any single founder's accumulated human capital. This refines the theoretical framework. It does not reject it. The signalling mechanism holds. The object being signalled is team structure, not founder pedigree.

The cohort results offer partial support for the macroeconomic moderation argument. Firms founded in 2023–26 are significantly more likely to secure financing than the 2009–12 base cohort ($\beta = .633$, $p = .035$). The intermediate cohorts show no significant difference, however.

This limits how far the results can be read as tracing a full hot-market and cold-market cycle, of the kind described by Nanda and Rhodes-Kropf (2013). The cross-sectional design cannot fully separate a genuine cohort effect from a broader market-cycle effect. We come back to this limitation in Section 6.6.

5.8.3 VC capital as an accelerator of polarised outcomes

Chapter 2 argued that VC capital does more than facilitate survival. It accelerates a venture towards extreme outcomes: a cash-out through acquisition or IPO, or a flame-out through shutdown (Arora and Nandkumar, 2009). The results only partly support this view.

The *financing-propensity* variable shows *no significant association* with overall survival in the unified Cox model. This null result could reflect two things. It could hide two offsetting effects, one pushing towards shutdown and one towards acquisition. Or it could simply be a genuine null. Section 5.7.3 settles this question.

The competing-risks results are clear. Financing propensity is not significant in either cause-specific model. It has no effect on shutdown ($p = .877$). It has no effect on acquisition ($p = .370$). This is a real null result, not a cancelled-out effect.

This finding does not confirm the accelerant framework proposed in Chapter 2. Financing propensity does not push firms towards either pole, cash-out or flame-out. In this sample, at least, it does not drive the timing or the mode of exit at all.

The sector results tell a different story, and here the evidence is stronger. As shown in Section 5.7.3, the elevated hazard recorded for Core cybersecurity firms in the unified model sits entirely within the acquisition model. It is absent from the shutdown model. Core cybersecurity firms face a significantly higher hazard of acquisition, around 78% higher than the base category. They show no distinguishable hazard of shutdown.

This pattern is consistent with the acquirer structure discussed in Chapter 3. Large platform acquirers and telecommunications carriers actively absorb frontier cybersecurity technology through acquisition. It also lends indirect support to the competence-destroying discontinuity framework (Tushman and Anderson, 1986). A segment exposed to rapid technological turnover appears more prone to acquisition for its capabilities than to outright failure.

CHAPTER 6: CONCLUSION

6.1 Summary

Chapter 2 of this thesis we outlined three theoretical strands. Founder background was framed as a signal to investors. Firm age was framed as the main driver of survival, through the Valley of Death. Venture capital was framed as an ‘accelerator’, pushing firms towards polarised outcomes, either a successful cash-out or a rapid failure. Each strand led to *three complementary research questions* which were then tested. The empirical results, read together, have told tell a more structural story than the literature predicts.

RQ1 asked what drives access to equity financing. Team composition emerges as the strongest predictor. Individual founder credentials carry no independent weight once team structure is accounted for: neither professional background nor educational background. Investors appear to evaluate the organisational structure of startups, not individual pedigree. Sector plays a more selective role. Firms in IT Services & Professional Consulting are significantly less likely to secure financing than the base category ($p = .025$), while founding cohort shows an effect only for the most recent period, 2023–26 ($p = .035$).

RQ2 asked what drives survival, conditional on financing. Firm age plays almost no role. The hazard of exit rises steadily across the observation window, without the early peak the Valley of Death literature predicts. Size and sector carry the explanatory weight instead. The hazard of exit rises monotonically with firm size, from 62.5% higher for firms with 11–200 employees to almost six times higher for firms above 1,000 employees. The Core cybersecurity cluster also shows a significantly higher hazard than the base category.

RQ3 asked how firms exit, conditional on financing. This is where the two-stage design pays off. The competing-risks results show that the higher hazard for larger firms in RQ2 is almost entirely driven by the hazard of acquisition, not of failure. The acquisition-model coefficient rises with size, while the shutdown-model coefficient shows no comparable pattern at the top of the size distribution.

Financing propensity, by contrast, shows no significant association with either shutdown or acquisition. It does not accelerate a firm towards one outcome or the other. It does not appear to shape the timing of exit at all, once shutdown and acquisition are treated as distinct events.

Three conclusions follow from this evidence.

- I. First, structure beats biography throughout. Team composition matters more than who the individual founders are. Firm size and sector specificity matter more than how long a firm has existed. In each case, a structural characteristic outperforms an individual or temporal one.
- II. Second, the ‘accelerator framework’ does not hold in this sample. Financing propensity does not push firms towards a faster cash-out or a faster failure. This is a genuine departure from the theory set out in Chapter 2, not merely a weak confirmation of it.
- III. Third, exit through acquisition is best read as a market outcome, not a failure. Larger firms, and firms in the Core cybersecurity cluster specifically, exit predominantly through acquisition. Outright failure stays concentrated among smaller, medium-sized, and lower-revenue firms. Vulnerability in this sector is therefore structural. It is shaped by size and sector positioning, not by firm age, founder background, or the propensity to raise capital.

6.2 Implications

These results carry several practical implications for investors, founders, and policy makers.

For investors. The findings suggest a specific shift in how risk should be assessed in this sector. Firm age is a weak signal of exit risk: the hazard of exit rises steadily across the observation window, without the early peak that standard startup risk models assume. Firm size and sector positioning are stronger signals. A firm's scaling stage, rather than its calendar age, should weigh more heavily in due diligence and portfolio monitoring. The size effect also carries a more specific

message. Most of the elevated hazard for larger firms reflects acquisition, not failure, so a rising hazard at this stage should not automatically be read as rising risk in the ordinary sense. For many firms, it signals proximity to a liquidity event rather than distress. A real vulnerability window does remain, however, concentrated in firms moving from small to medium size, where structural costs rise sharply. This is the stage at which monitoring should intensify.

For founders. The results point to a specific, actionable signal in the financing process. Team composition is the strongest predictor of securing equity financing, well ahead of any individual founder's professional or educational credentials. A founder building a solo venture faces a materially harder financing path than one who builds a team, independently of personal pedigree. For founders in the Core cybersecurity segment specifically, the results also carry a strategic implication: firms in this cluster face a significantly higher likelihood of acquisition rather than shutdown, which reframes acquisition readiness, rather than pure survival, as a relevant long-run objective for many ventures in this space.

For policy makers. The flat baseline hazard is the starting point. Risk in this sample is not concentrated in the earliest years. Many public support schemes assume otherwise. They target the first three to five years of a firm's life, often through tax relief or grant eligibility. This calibration may be poorly suited to cybersecurity. An eight-year-old firm undergoing rapid scaling can carry more exit risk than a firm in its first year. Eligibility criteria should therefore weigh firm size and sector cluster more heavily. Firm age alone is a weak guide.

The size effect points to a related concern. Most of the elevated hazard for larger firms reflects acquisition, not failure. A real vulnerability window remains, though, and it is concentrated in a specific place: firms in the 201–1,000 employee band carry a significantly higher hazard of shutdown than any other size category. This is the scaling phase itself, when structural costs rise sharply and the firm has outgrown early-stage support but has not yet reached the stability of the largest players. Liquidity support targeted at this transition would address a risk that current early-stage schemes largely miss.

The results for the Core cybersecurity segment signal a further issue. Its high exit rate is driven by acquisitions, not by operational failure. Firms in this cluster face a hazard of acquisition around 78% higher than the base category, with no comparable increase in shutdown risk. This is not, in itself, a bad outcome. Frontier technology is being absorbed by larger players. But it raises a longer-term concern about market concentration. If the most innovative segment of the industry is being consistently acquired by a small number of large players, competition policy has a role to play. Sector-specific antitrust monitoring of cybersecurity M&A would be a reasonable response, alongside targeted support for start-ups whose independence carries strategic value, for instance in areas touching data sovereignty or national security.

The financing-propensity results carry a different lesson than the accelerator framework in Chapter 2 might suggest. Financing propensity shows no significant association with either the timing or the mode of exit, once shutdown and acquisition are treated separately. Equity financing does not appear to push firms towards a faster cash-out, nor towards a faster failure, at least on the evidence of this sample. This matters for public co-investment schemes that treat venture capital access as a marker of reduced risk. Securing external funding is not, by itself, a signal that a firm is on a safer path. Public venture funds and co-investment vehicles should therefore weigh structural characteristics, such as size and sector, more heavily than funding status when assessing a firm's risk profile and should pair any funding with operational support during the scaling phase rather than treating capital access as sufficient on its own.

Regulation adds one further dimension, already visible in the sector's history. The NERC-CIP standard, discussed in Chapter 3, triggered a concentrated wave of acquisitions when it came into force. New regulatory frameworks in cybersecurity, such as the EU's NIS2 directive, could plausibly have a similar effect. Regulators should anticipate this. A compliance deadline does not only raise technical standards. It can also accelerate market consolidation, as smaller firms are absorbed by larger, already-compliant acquirers rather than adapting independently.

6.3 Limitations

This thesis is subject to several limitations, which can be used to qualify the conclusions drawn above.

The dataset is drawn from Crunchbase, a partial data source which is partly based on information that are self-reported by the companies. It tends to over-represent firms with visible, traceable funding events. It under-represents informal shutdowns, quiet closures, and exits that go unreported. The true incidence of shutdown in this sector may therefore be higher than the 3.6% recorded in the sample. This is a common limitation in startup survival research, not one specific to this thesis, but it should be borne in mind when interpreting the exit rates presented above.

The sample size is 1,227 firms. This is adequate for the logit and Cox specifications used. It is more limiting for the competing-risks extension, where only 159 firms recorded an exit, split further into 44 shutdowns and 113 acquisitions. Several size and revenue categories contain zero shutdown events, producing numerically degenerate coefficients in the cause-specific shutdown model. These coefficients were flagged and excluded from substantive interpretation, but the underlying constraint remains: cause-specific estimates rest on a comparatively small number of exit events.

Founder background shows limited variation within this sample. Professional and educational backgrounds are relatively homogeneous across firms. Most founders share a similar technical profile. This homogeneity may explain why professional and educational background show no significant effect on financing probability, a result that sits in tension with parts of the human-capital literature reviewed in Chapter 2. A more heterogeneous sample, spanning a wider range of founder profiles, might reveal effects that this dataset cannot detect. The non-result should therefore be read as a feature of this sample, not necessarily as a rejection of the underlying theory.

The competing-risks framework captures only two exit routes: shutdown and acquisition. The dataset contains no IPOs and no publicly listed firms at the time yhr data was collected. All sample firms are privately held. This has two consequences. First, the "cash-out via IPO" pathway discussed in Chapter 2 cannot be tested directly; acquisition stands as the only observable proxy for a successful liquidity event. Second, all sample firms are based in Europe. Financing conditions, acquirer behaviour, and regulatory pressure all differ across regions, and the patterns identified here may not hold in the same form in the United States or in Asia, where venture capital markets and M&A activity follow different dynamics.

The time horizon of the sample, 2009 to 2024, spans several distinct financing cycles. Hot and cold markets are only partially captured through cohort fixed effects. These effects cannot fully separate a true cohort difference from a broader market-cycle effect operating at the same time. A longer, more granular time series would allow these two sources of variation to be distinguished more precisely.

Finally, the instrumental approach used in Stage 2, the predicted financing propensity, addresses part of the endogeneity concern raised by using financing as a survival predictor. However, it does not eliminate it entirely. Unobserved firm-level characteristics may still influence both financing access and exit outcomes jointly. The results should therefore be read as strongly suggestive rather than as fully causal evidence, in keeping with the hedged, non-causal language adopted throughout this thesis.

Taken together, these limitations bear with particular weight on the answer to RQ3. The null result found for financing propensity, robust across both the unified and the cause-specific models, is unlikely to be an artefact of the centring procedure or of a single specification, since it holds consistently across every version tested. It is more plausibly connected to the sample-size and event-count constraints described above: with only 44 shutdowns and 113 acquisitions to estimate from, the cause-specific models may lack the statistical power to detect a genuinely modest effect of financing propensity, even where one exists. Moreover, the absence of IPO exits and the exclusively European scope of the sample further limit how far this null

result can be generalised to other exit routes or other financing environments. The finding should therefore be read as a genuine and carefully verified pattern within this dataset, rather than as a definitive rejection of the accelerant hypothesis at the level of the broader cybersecurity sector.

BIBLIOGRAPHY

Agarwal, R., & Gort, M. (2002). Firm and product life cycles and firm survival. Department of Business Administration, University of Illinois, 184–190.

Allen, H. (2025). Interest rates, venture capital, and financial stability. *University of Illinois Law Review*, 2025(5).

Alunni, A. (2020). La importancia de la inversión de riesgo corporativo para superar el valle de la muerte en la transferencia tecnológica universitaria. *Revista de Fomento Social*, 1, 139–156.

Amit, R., Muller, E., & Cockburn, I. (1995). Opportunity costs and entrepreneurial activity. *Journal of Business Venturing*, 10(2), 95–106.

Arora, Nandkumar (2009). Cash-out or flame-out! Opportunity cost and entrepreneurial strategy: Theory, and evidence from the information security industry.

Belenzon, S., Chatterji, A., & Daley, B. (2017). Eponymous entrepreneurs. *American Economic Review*, 107(6), 1638–1655.

Belenzon, S., Chatterji, A., & Daley, B. (2018). Choosing between growth and glory.

Brüderl, J., & Schüssler, R. (1990). Organizational mortality: The liabilities of newness and adolescence. *Administrative Science Quarterly*, 35(3), 530–547.

Esteve-Pérez, S., & Mañez-Castillejo, J. A. (2008). The resource-based theory of the firm and firm survival. *Small Business Economics*, 30, 231–249.

Fine, J.P. and Gray, R.J. (1999) 'A proportional hazards model for the subdistribution of a competing risk', *Journal of the American Statistical Association*, 94(446), pp. 496–509.

Franco, M., & Haase, H. (2010). Failure factors in small and medium-sized enterprises: Qualitative study from an attributional perspective. *International Entrepreneurship and Management Journal*, 6, 503–521.

Greenberg, J., & Mollick, E. (2018). *Sole survivors: Solo ventures versus founding teams*. SSRN. <https://doi.org/10.2139/ssrn.3107898>

Hejazi, S. R., Seyyedamiri, N., Shokouhi, M., Bidhandi & Rasoulilian, P. (2024). Identifying factors affecting startups survival: A systematic literature review.

Honjo, Y., Iwaki, Y., & Kato, M. (2024). Outside or inside the firm? The impact of debt financing on the exit routes of start-up firms.

Hsu, D. H. (2007). Experienced entrepreneurial founders, organizational capital, and venture capital funding. *Research Policy*, 36(5), 722–741.

Josefy, M. A., Harrison, J. S., Sirmon, D. G., & Carnes, C. (2017). Living and dying: Synthesizing the literature on firm survival and failure across stages of development. *Academy of Management Annals*, 11(2), 770–799.

Ko, E.-J., & McKelvie, A. (2018). Signaling for more money: The roles of founders' human capital and investor prominence in resource acquisition across different stages of firm development. *Journal of Business Venturing*, 33(4), 438–454.

Lyles, M. A., Saxton, T., & Watson, K. (2004). Venture survival in a transitional economy. *Journal of Management*, 30(3), 351–375.

Ma, Y., & Zimmermann, K. (2023). *Monetary policy and innovation* (NBER Working Paper No. 31698). National Bureau of Economic Research.

Moiseienko, T., & Kiva, A. (2022). Cybersecurity startup investments. *CEUR Workshop Proceedings*, 3188, 244–250.

Nanda, R., & Rhodes-Kropf, M. (2013). Investment cycles and startup innovation. *Journal of Financial Economics*, 110(2), 403–418. <https://doi.org/10.1016/j.jfineco.2013.07.001>

OCDE. (2015). *Promoviendo el desarrollo de sistemas locales de innovación.*

Paik, Y. (2014). Serial entrepreneurs and venture survival: Evidence from US venture-capital-financed semiconductor firms. *Strategic Entrepreneurship Journal*, 8, 254–268.

Simpeh, K. N. (2011). Entrepreneurship theories and empirical research: A summary review of the literature. *European Journal of Business and Management*, 3(6), 1–8.

Stinchcombe, A. L. (1965). Social structure and organizations. In J. G. March (Ed.), *Handbook of organizations* (pp. 142–193). Rand McNally.

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

Zapata-Molina, C., Montes-Hincapié, J. M., Londoño-Arias, J. A., & Baier-Fuentes, H. (2022). The valley of death of start-ups: A systematic literature review.

Zhang, C. (2016). *A study on cybersecurity start-ups: A financial approach to analyze industry trends, entrepreneurship ecosystems and start-up exits* [Master's thesis, Massachusetts Institute of Technology]. MIT Sloan School of Management.

Zygierewicz, A. (2020). *Directive on security of network and information systems (NIS Directive): EU cybersecurity policy* (Implementation Appraisal, PE 654.198). European Parliamentary Research Service (EPRS), European Parliament.