LUISS

Master Degree Program in Management
Major in International Management

Course of Advanced Corporate Finance

Predicting Strategic Retrenchment: A Volatility-Based Analysis of IPO Withdrawal

Prof. Fabrizio Core

SUPERVISOR

Prof. Rosella Santella

CO-SUPERVISOR

Luca Giordano

CANDIDATE

Academic Year 2025/2026

Predicting Strategic Retrenchment:

A Volatility-Based Analysis of IPO Withdrawal

Index

1 Introduction

1.1 IPOs, External Financing, and Market Conditions

1.1.1 Role of external financing in corporate finance

Corporate finance is centrally concerned with how firms raise and allocate financial resources in order to maximize firm value over time. Alongside investment and payout decisions, financing choices determine the capacity of firms to undertake profitable projects, sustain long-term growth, and respond to changing economic environments. While internal funds generated through retained earnings represent a natural and relatively inexpensive source of capital, they are often insufficient to finance large-scale investments or periods of rapid expansion. In such cases, firms must rely on external financing, typically through debt or equity issuance. The selection among alternative sources of external capital constitutes one of the most fundamental problems in corporate finance, as each financing instrument entails distinct implications for risk allocation, ownership structure, control rights, and the overall cost of capital.

The theoretical benchmark for analysing financing decisions is provided by the irrelevance proposition of Modigliani and Miller (1958)[1], which states that in frictionless capital markets characterized by perfect information, no taxes, and no transaction costs, the value of the firm is independent of its capital structure. Under these conditions, firms can raise external funds costlessly, and investors can fully diversify risk, rendering financing choices neutral with respect to value creation. Although this framework offers a powerful analytical starting point, it is widely recognized that its assumptions are highly restrictive. Once real-world frictions are introduced, financing decisions become economically meaningful and can exert a significant influence on firm behaviour and valuation. Consequently, modern corporate finance research focuses on identifying the mechanisms through which external financing affects firm value and on understanding how firms optimally choose among financing alternatives in imperfect markets.

A central insight of this literature is that external equity financing is particularly sensitive to informational frictions. Myers and Majluf (1984)[2] demonstrate that when managers possess superior information about firm value relative to outside investors, the issuance of new equity may be interpreted as a negative signal regarding the firm's prospects. Anticipating this adverse selection problem, rational investors demand a discount to compensate for the risk of overvaluation, thereby increasing the cost of equity financing. As a result, firms may refrain from issuing equity even when profitable investment opportunities are available, preferring internal funds or debt financing instead. This pecking order behaviour underscores that equity issuance is not a purely mechanical funding choice, but a strategic decision shaped by information asymmetries, investor beliefs, and market perceptions.

Beyond informational considerations, corporate finance theory emphasizes that financing decisions are inherently forward-looking and made under uncertainty. When firms raise external capital, they commit to future obligations and expose themselves to uncertain future states of the world. The

[1] Modigliani, F., & Miller, M. H. (1958). The cost of capital, corporation finance, and the theory of investment. *American Economic Review, 48*, 261–297.

[2] Myers, S. C., & Majluf, N. S. (1984). *Corporate financing and investment decisions when firms have information that investors do not have* (NBER Working Paper No. 1396). National Bureau of Economic Research.

benefits of financing depend not only on expected cash flows, but also on the distribution of possible outcomes and the risks associated with them. In the context of equity issuance, uncertainty affects both the feasibility and the effective cost of raising capital, as investors require compensation for bearing risk and for committing funds in environments characterized by heightened uncertainty. From the firm's perspective, this implies that financing decisions cannot be evaluated solely on the basis of expected valuations but must account for the risks associated with future market conditions and investor responses.

The role of uncertainty in financing decisions has been further clarified through the application of real options theory to corporate finance. Building on the work of Bernanke (1983)[3] and Dixit and Pindyck (1994)[4], this strand of the literature argues that when investment or financing decisions are at least partially irreversible, uncertainty increases the value of waiting. Raising external capital often involves substantial sunk costs, including underwriting fees, disclosure expenses, legal costs, and potential reputational consequences. Once these costs are incurred, they cannot be fully recovered if conditions deteriorate. As a result, firms may rationally delay or abandon financing plans when uncertainty is elevated, even if immediate financing appears attractive based on expected values alone. This real-options perspective highlights that the timing of financing decisions is itself a strategic variable and that firms optimally balance the benefits of early financing against the option value of waiting for more favorable conditions.

Within this broader corporate finance framework, equity financing through public markets occupies a particularly important position. Going public through an Initial Public Offering (IPO) represents a major transformation in a firm's financing structure, governance arrangements, and relationship with capital markets. Public equity markets provide access to large pools of capital, enhance liquidity for existing shareholders, and facilitate future financing through seasoned equity offerings or debt issuance. At the same time, public listing subjects firms to continuous market valuation, regulatory oversight, and disclosure requirements, which can influence managerial incentives and strategic flexibility. The decision to raise capital through an IPO therefore reflects a complex trade-off between financing capacity, cost of capital, control considerations, and exposure to market forces.

Empirical evidence consistently shows that firms do not access public equity markets randomly over time. Instead, equity issuance exhibits pronounced cyclicality, with periods of intense issuance activity followed by prolonged declines. Baker and Wurgler (2002)[5] interpret this pattern as evidence that firms actively time their financing decisions in response to market conditions. According to this view, managers are more likely to issue equity when market valuations are perceived to be high and when investor demand for new issues is strong, while postponing issuance during unfavourable periods. Although the precise motivations underlying market timing remain the subject of ongoing debate, the empirical regularity itself underscores a fundamental insight of

[3] Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics, 98*(1), 85–106.

[4] Dixit, A., & Pindyck, R. (1994). *Investment under uncertainty.* Princeton University Press.

[5] Baker, M., & Wurgler, J. (2002). Market timing and capital structure. *The Journal of Finance, 57*(1), 1–32.

corporate finance: financing decisions are shaped not only by firm-specific characteristics and funding needs, but also by broader market conditions.

This cyclical behaviour is particularly evident in the market for initial public offerings. Decades of research document that IPO activity tends to surge during so-called "hot markets" and decline sharply during downturns (Loughran, Ritter, & Rydqvist, 1994[6]; Ritter & Welch, 2002[7]). Firms appear to strategically choose when to go public in order to maximize the proceeds of the offering and to improve the likelihood of a successful outcome. Historical episodes illustrate this pattern clearly, with IPO volumes peaking during periods of strong equity markets and contracting during episodes of financial distress, such as the global financial crisis of 2008. These dynamics suggest that the decision to go public is deeply intertwined with the state of financial markets and cannot be understood independently of prevailing market conditions.

Importantly, the sensitivity of IPO activity to market conditions has implications not only for the volume of offerings, but also for their outcomes. Research has shown that IPO pricing, underpricing, and long-run performance vary systematically across issuance cycles and regulatory regimes. High-issuance periods often coincide with lower average underpricing, while firms that go public at market peaks may experience weaker post-IPO performance, consistent with the exploitation of temporary windows of opportunity (Ritter, 1991[8]; Loughran & Ritter, 1995[9]). These findings reinforce the view that equity issuance decisions are influenced by managerial assessments of market conditions and investor sentiment, rather than being driven solely by long-term fundamentals.

Taken together, the corporate finance literature establishes that external financing decisions, particularly those involving public equity, are complex, strategic, and highly sensitive to uncertainty and market conditions. Firms must weigh internal funding capacity against the costs and risks associated with accessing external capital, taking into account information asymmetries, irreversibility, and timing considerations. The decision to go public represents a culmination of these trade-offs, as it exposes firms to market valuation and investor behaviour in a highly visible and consequential manner.

1.1.2 IPOs as market-dependent transactions: Equity Financing and the Decision to Access Public Markets

Building on the role of external financing in corporate finance, an IPO represents a form of external equity financing, and specifically, is the process by which a private firm offers shares to public investors for the first time, transforming the firm into a publicly traded company. This transition is strategic: it opens access to a broad pool of capital but also introduces new obligations and risks.

---

[6] Loughran, T., Ritter, J. R., & Rydqvist, K. (1994). Initial public offerings: International insights. *Pacific-Basin Finance Journal, 2*, 165–199.

[7] Ritter, J. R., & Welch, I. (2002). *A review of IPO activity, pricing and allocations* (Yale ICF Working Paper No. 02-01). Yale International Center for Finance.

[8] Ritter, J. R. (1991). The long-run performance of initial public offerings. The Journal of Finance, 46, 3–27.

[9] Loughran, T., & Ritter, J. R. (1995). The new issue puzzle. *The Journal of Finance, 50*, 23–51.

Corporate finance theory traditionally emphasizes that external equity is a costly financing source due to information asymmetries between managers and investors. According to the pecking order theory (Myers & Majluf, 1984), firms prefer to use internal funds first, then debt, and only turn to equity as a last resort because issuing equity can signal to outsiders that the firm's stock may be overvalued. In this view, managers reluctant to dilute ownership will avoid equity issuance unless cheaper sources (retained earnings or additional debt) are insufficient. Myers (1984) further argued that firms might hoard "financial slack" to avoid stock issuance, since equity offerings face adverse selection costs, outside investors fear managers issue stock when it is overpriced (Brau & Fawcett, 2006[10]). Despite this inherent cost, when a firm's growth opportunities outstrip its internally available funds and debt capacity, equity financing becomes necessary. The IPO is one avenue to raise such equity on a large scale, marking a firm's entry into public capital markets.

1.1.2.1 Choosing to Go Public vs. Staying Private

The decision to go public can be viewed as a trade-off: firms must weigh the benefits of raising capital from public markets against the costs of broader ownership and disclosure. Chemmanur and Fulghieri's (1999)[11] theoretical model casts the IPO decision in terms of life-cycle and information trade-offs. At some stage of development, a company may choose between private equity financing (e.g. venture capital or private placements) and a public offering. In their model, the entrepreneur has private information about the firm's value, and outside investors can expend resources to evaluate the firm. Going public means selling shares to many dispersed investors, which can lead to duplication of information production as each investor seeks to evaluate the firm. By contrast, financing from a single private investor (such as a venture capitalist) concentrates information gathering but comes at the cost of a high risk premium demanded by that investor. Thus, the firm's trade-off is between the IPO's broader investor base (with inevitably higher disclosure and information-production by many small investors) and the more onerous terms of private financing (where a large investor demands significant returns and control rights). Chemmanur and Fulghieri argue that as firms become larger and more mature, the balance tilts toward public markets: by going public, a company can avoid giving an unduly large stake or control to a single investor, even though it must then bear the costs of regulatory disclosure and reduced privacy. This perspective aligns with a life-cycle view of financing, where young firms rely on private equity, but successful growth eventually pushes them to "graduate" to public equity markets in order to obtain larger, more liquid capital infusions.

Empirical evidence supports several aspects of this life-cycle progression. Pagano, Panetta, and Zingales (1998)[12] find that in Italy the likelihood of an IPO is higher for larger firms, and those in industries with high market valuations, consistent with the notion that companies go public when

---

[10] Brau, J. C., & Fawcett, S. E. (2006). Initial public offerings: An analysis of theory and practice. *The Journal of Finance, 61*(1), 399–436.

[11] Chemmanur, T. J., & Fulghieri, P. (1999). A theory of the going-public decision. *The Review of Financial Studies, 12*(2), 249–279.

[12] Pagano, M., Panetta, F., & Zingales, L. (1998). Why do companies go public? An empirical analysis. *The Journal of Finance, 53*(1), 27–64.

they have achieved a certain scale and operating environment to justify the fixed costs of listing. Notably, Pagano et al. report that firms often go public not only to finance new investment, but also to rebalance their balance sheets and allow early investors to diversify. In their sample, IPO proceeds were largely used to pay down debt (deleveraging) rather than to fund fresh expansion, and many IPOs involved controlling shareholders selling a portion of their equity stake. This suggests that firms use IPOs as an opportunity for insiders to liquidate or diversify holdings, especially after a period of rapid private growth. Going public was also associated with a reduction in the cost of bank credit for these firms, implying that a stock market listing can enhance a company's creditworthiness and bargaining power with lenders. Furthermore, Pagano et al. document a higher turnover of corporate control after IPO; newly public firms saw more frequent changes in ownership or takeovers than comparable private firms. In other words, public status increases a firm's exposure to the market for corporate control, which can be a double-edged sword: it facilitates acquisitions and alliances (using stock as currency) but also makes the firm more contestable (easier to be acquired or subject to activist investors). These findings highlight that the motives for going public can extend beyond simply raising cash for investment; they include balance-sheet restructuring, investor exits, and positioning the firm for strategic opportunities.

1.1.2.2 Key Motivations and Benefits of IPOs

The primary impetus for an IPO is indeed to raise capital for growth. By selling shares to the public, a company can obtain a large injection of equity funding to support expansion projects, R&D, or other investments that private financing might not sufficiently cover. An IPO also broadens the investor base dramatically: instead of relying on a handful of private investors, the firm can tap potentially thousands of institutional and retail shareholders domestically and abroad. This broader ownership brings several advantages. First, it often creates a market-valued currency: publicly traded stock. Once listed, a firm's shares have an observable market price and can be readily traded, which enables their use in non-cash acquisitions and mergers. Brau and Fawcett (2006), who surveyed CFOs of firms that had gone public, found that funding growth opportunities was the top-ranked motive for going public, with executives specifically highlighting the ability to use public stock to finance acquisitions as a critical advantage. This underscores that IPOs are not only about raising cash now, but also about enabling strategic flexibility: a listed company can grow by acquiring other firms and paying with its own shares. Relatedly, having a publicly traded stock establishes a market price for the company's equity, which provides an objective benchmark of firm value. CFOs in Brau and Fawcett's study cited the desire to obtain a market valuation as another major reason to go public. A market valuation is useful for multiple reasons; it informs corporate decisions, allows management to tie stock-based compensation to a transparent metric, and can signal the firm's worth to potential partners or acquirers.

Being public can also enhance a company's visibility and credibility. Listing on a stock exchange often increases media and analyst coverage of the firm, raising its profile. The added transparency and scrutiny can improve the firm's reputation with customers, suppliers, and employees. For example, the requirement to publish audited financial statements and comply with exchange regulations signals that the company meets higher governance standards. Classic corporate finance texts note that a public listing imposes discipline on management: the stock price serves as an ongoing performance measure and any missteps by management are quickly reflected in a falling share price. This market discipline, along with regulations protecting minority shareholders, may help curb agency problems by aligning managers' incentives more closely with shareholders (e.g. through stock option plans). Additionally, public companies can more readily attract top talent by offering liquid equity (stock or stock options) as part of compensation, an incentive mechanism much less feasible in private firms. Finally, going public can provide liquidity for existing shareholders. Founders, early employees, or venture capital investors often cannot easily sell their stake in a private company; an IPO creates a public market for the shares, allowing these insiders to gradually cash out or diversify their personal wealth. Indeed, allowing principals to diversify or exit is an acknowledged motive in the IPO decision. While CEOs do consider this, it typically ranks somewhat below the growth-related motives in importance. Nevertheless, many IPOs include a "secondary" component (shares sold by existing owners) alongside the primary issuance of new shares, indicating that insiders are taking the opportunity to reap liquidity. Pagano et al. (1998) explicitly found that many Italian IPOs coincided with equity sales by controlling shareholders, consistent with the diversification motive.

1.1.2.3 How IPOs Differ from Private Equity Financing

It is instructive to contrast IPOs with alternative equity financing routes such as private equity or venture capital. One key difference lies in the ownership structure and control. A venture capital firm or private equity fund that invests in a private company typically takes a significant ownership stake and often demands board seats, veto rights, and covenants to protect its investment. This concentrated ownership can benefit the firm through active monitoring and expertise, but it also means a loss of autonomy for the founders. By going public, a company instead disperses ownership among many shareholders, none of whom individually may have substantial control (unless dual-class shares or a controlling shareholder structure is maintained). The dispersed ownership of public equity can reduce direct interference in management; public shareholders generally vote on major issues but delegate day-to-day control to professional managers. However, dispersed shareholders are also less able to coordinate, which can lead to different governance challenges (e.g. the classic separation of ownership and control). Another difference is the price-setting mechanism and information environment. In a private financing round, the firm's value is negotiated behind closed doors with the investor, who gains access to internal data (and often conducts due diligence) before committing funds. In an IPO, pricing is determined via a market process; underwriters gather demand from many investors (book-building) to set a public offer price, and once trading begins, the stock price is continuously determined by the market. This market-driven pricing provides transparency but is subject to volatility and investor sentiment in a way that private deals are not. Because public investors have less access to private information than, say, a venture capitalist would, information asymmetry is a central concern in IPOs. Firms address this by producing a detailed prospectus and financial disclosures, and by hiring reputable investment banks as underwriters to market the offering. The underwriter's role, as described by Booth and Smith (1986)[13], is partly to certify the firm's value to the market. By backing the IPO, especially if the underwriter is well-known and prestigious, they send a signal that a credible intermediary has

---

[13] Booth, J. R., & Smith, R. L. (1986). Capital raising, underwriting and the certification hypothesis. *Journal of Financial Economics, 15*, 261–281.

vetted the company's quality. This certification hypothesis posits that underwriters use their reputational capital to assure investors that the offering price reflects all available information (even potential adverse information) about the firm. In effect, the underwriter helps bridge the information gap, reducing the risk perceived by outside investors. This is an important distinction between IPOs and private equity: in private deals, the investor's own due diligence and contractual protections mitigate information problems, whereas in an IPO the assurance is provided via certification, mandatory disclosures, and regulatory oversight.

1.1.2.4 Costs and Drawbacks of Going Public

Against these benefits, a firm must consider the costs of accessing public markets. One well-documented cost is IPO underpricing; the tendency for the IPO offer price to be set below the price at which the stock trades once it is on the market. Underpricing generates a first-day return for new investors but represents "money left on the table" for the issuing firm. In their comprehensive survey, Ritter and Welch (2002)[14] report that, in the U.S., IPO shares between 1980–2001 closed nearly 19% higher on average on the first trading day compared to the IPO offer price.
In other words, companies could theoretically have sold shares at a significantly higher price; the wealth difference accrues to initial investors at the expense of the original shareholders. Underpricing is thus an indirect cost of issuance, often justified as a way to compensate investors for information uncertainties or to ensure a successful sale, but nonetheless a real cost to founders and pre-IPO owners. In addition, companies incur substantial direct costs in the IPO process there are additional expenses for legal counsel, accounting, regulatory filing, and marketing the offering. These fees make IPOs one of the most expensive financing methods. Thereafter, as a public company, the firm faces ongoing costs of compliance and disclosure: annual and quarterly financial reporting, audit requirements, adherence to stock exchange rules and corporate governance codes, and the need to communicate with a large shareholder base. Preparing and disseminating this information is costly and may force the firm to reveal competitively sensitive data that it would prefer to keep private (e.g. detailed financials or strategic plans that rivals could learn from). Another major consideration is the loss of control and privacy. Going public dilutes the ownership stake of the original owners, potentially reducing their control over corporate decisions. In many cases, founders can retain control by holding a majority of shares or creating dual-class share structures, but such arrangements may not always be possible or desired. With a broad set of shareholders, management becomes accountable to public investors and may come under pressure from activist shareholders or face hostile takeover attempts if the firm underperforms. Brau and Fawcett (2006) find that the primary reason executives cite for staying private is to maintain decision-making control over the company. Privately held firms can make strategic decisions with a long-term horizon away from the public eye, whereas public firms must contend with market expectations and fluctuations. The need to meet quarterly earnings targets or maintain the stock price can introduce short-term pressures that might conflict with long-term investment plans.

---

[14] Ritter, J. R., & Welch, I. (2002). *A review of IPO activity, pricing and allocations* (Yale ICF Working Paper No. 02-01). Yale International Center for Finance.

Moreover, public shareholders expect a say in significant matters (through board elections, proxy votes, etc.), and regulatory rules often require certain governance structures (e.g. independent directors, audit committees) once a firm is listed, further checking the absolute authority of founding owners.
Finally, IPOs carry reputational risks. A failed IPO, for instance, if the offering has to be withdrawn due to insufficient investor demand, or if the stock performs poorly right after listing, can tarnish a company's reputation in the markets. Managers often worry that an unsuccessful offering might signal weakness, making it harder to raise capital in the future or damaging the firm's image with customers and employees. Even after a successful IPO, the firm's leadership is under the continuous scrutiny of analysts and investors; any missteps can result in stock price declines and negative publicity. Thus, going public exposes the firm to the discipline of the market, which is beneficial when performance is strong but unforgiving when challenges arise. In weighing an IPO, companies must therefore consider not only the upfront benefits (capital, liquidity, growth opportunities) but also the long-run commitments they are making to greater transparency, shared governance, and market-driven evaluation of their performance.

1.1.3 Introduction of market conditions and uncertainty

Having established why firms may choose to raise equity through public markets, we now consider when they choose to do so. The timing of an initial public offering is a critical strategic decision in corporate financing, one that extends beyond a firm's internal needs and characteristics. A robust finding in finance is that IPO activity is highly cyclical, tending to cluster in "hot" markets. In other words, IPOs often come in waves: certain periods witness a glut of offerings with buoyant first-day returns, while other "cold" periods see few new listings. This empirical reality of IPO waves implies that broader market conditions, not just a firm's own readiness, heavily influence the decision of *when* to go public.

Researchers first documented the "hot issue market" phenomenon in the 1970s. Ibbotson and Jaffe (1975)[15] observed that there are periods of IPO "euphoria" in which the number of offerings surges and average initial returns are unusually high. For instance, during the late 1990s tech boom, one of the hottest IPO markets on record, hundreds of companies went public amid exuberant investor sentiment, producing unprecedented total proceeds and large first-day price jumps. Such peaks eventually cool off; IPO volume falls during "cold" market windows when investor appetite for new equity issues diminishes. IPO cycles are a fundamental feature of equity markets that any theory of the going-public decision must account for. In practice, issuers and underwriters commonly refer to the existence of an "open" or "closed" IPO window depending on current market receptivity.

1.1.3.1 Market conditions and sentiment

The clustering of IPOs in hot markets suggests that firms try to time their offerings to coincide with favorable market conditions. In hot markets, stock prices are generally rising, investor sentiment is

[15] Ibbotson, R. G., & Jaffe, J. F. (1975). "Hot issue" markets. *The journal of finance*, *30*(4), 1027-1042.

bullish, and the cost of equity capital is perceived to be relatively low. In such environments, investors are optimistic and hungry for new investment opportunities, creating what issuers see as a *window of opportunity* to sell new shares at attractive valuations. By contrast, during market lulls or downturns, investors become more risk-averse and liquidity for new issues dries up; a discouraging scenario for firms contemplating an IPO. This underscores that IPO timing is not solely driven by firm-specific factors like growth prospects or financing needs; it is also shaped by external forces such as overall market sentiment and the availability of risk capital.
Academic research corroborates that both investor sentiment and the demand for capital contribute to IPO waves. Lowry (2003)[16] finds that periods of high IPO volume can be partly explained by optimistic investor sentiment, for instance, strong recent stock market returns or low yields driving investors toward equities, and partly by firms' aggregate demand for capital to fund new investments. In other words, companies tend to go public when equity market conditions are favorable (high valuations and upbeat sentiment) and when they have worthwhile projects that need financing. These "pull" and "push" factors are not mutually exclusive; a booming market often coincides with technological or economic developments that create new firms and growth opportunities, fuelling both the supply of IPO candidates and investor demand for shares.

1.1.3.2 Windows of opportunity and market timing theory

A large body of literature interprets IPO waves through the lens of market timing**.** The core idea is that corporate managers time equity issuances to exploit transient mispricing or valuation peaks. For example, Baker and Wurgler (2002) show that firms are far more likely to issue equity when their market valuations (e.g. market-to-book ratios) are high, and conversely hold back when valuations are low. In their words, capital structures at any point in time reflect the cumulative outcome of past attempts to time the equity market. In the IPO context, this means firms will prefer to go public when they believe their stock can fetch a relatively high price, typically in a bull market or a sector boom. Consistent with this view, Loughran, Ritter, and Rydqvist (1994) hypothesize that issuers take advantage of *"windows of opportunity"* opened by overoptimistic investors. During euphoric periods, investors may temporarily overvalue growth prospects or underplay risks, allowing even lesser-known or riskier firms to obtain capital on favorable terms. Firms rationally respond by accelerating IPO plans to "cash in" on strong market sentiment before the window closes.

It is important to note that going public during a hot market can influence IPO pricing and outcomes. Issuing in a frothy market often entails accepting higher underpricing (initial return) as a trade-off for a higher offering price and a greater chance of a fully subscribed issue. Indeed, hot issue markets are characterized by unusually high first-day returns on average, which indicates that IPOs in these periods leave more money on the table. One explanation is that underwriters set offer prices conservatively when investor demand is abundant, resulting in large pops, to ensure successful sales amid a glut of offerings. Another possibility is that firms going public in waves intentionally underprice more to signal their quality or to stand out in a crowded field, a theory related to IPO

---

[16] Lowry, M. B. (2003). Why does IPO volume fluctuate so much? *Journal of Financial Economics, 67*, 3–40.

timing as a signal in itself (as explored by recent models of IPO timing dynamics). Furthermore, there is evidence that IPOs launched in hot markets can experience different long-run performance than those from colder periods. Some studies find that companies going public in exuberant markets underperform in the long run, consistent with the idea that they were overvalued at issuance. However, distinguishing true market timing from mere "pseudo-timing" effects (where waves naturally coincide with market peaks without managers perfectly foreseeing the peak) has been a subject of debate. Overall, the consequences of timing reinforce its importance: picking the wrong moment to go public can mean a lower price, higher issuance costs, or poor aftermarket results for the issuing firm.

1.1.3.3 Rational theories of IPO waves

While sentiment and opportunism clearly play a role, not all scholars attribute IPO waves purely to irrational market timing. An alternative perspective comes from rational, information-based models. Pástor and Veronesi (2005)[17] develop a theory of "rational IPO waves" in which waves emerge as optimal responses to changing market conditions, even without assuming investor irrationality. In their model, the number of firms going public fluctuates because firms wait for certain market indicators to improve before going public. Specifically, declines in the expected market return (which lower the cost of equity) or increases in expected aggregate profitabilit**y** make conditions more favourable for IPOs, triggering more firms to file. Waves can also be driven by the resolution of uncertainty: if there is high prior uncertainty about a new industry or technology, firms may delay going public until some uncertainty is resolved or until a peer firm's successful IPO provides a valuation benchmark. Once that happens, a cluster of IPOs follows as firms capitalize on the now-improved information environment. This can give the appearance of a sudden "wave". Pástor and Veronesi's empirical tests support aspects of this theory, for example, they find IPO waves tend to be preceded by high stock market returns and often followed by lower returns, consistent with the notion that firms try to go public when conditions have been strong but may soften afterward. In sum, broader market conditions (expected returns, economic prospects, and information flows) can rationally drive the timing of IPOs, alongside the more opportunistic factors.
Crucially, adverse market conditions often deter IPO activity**.** Firms do not operate in a vacuum; if the market backdrop is unfavourable, even companies in need of capital may choose to postpone an offering. High stock market volatility is one such deterrent. Volatility represents uncertainty about market prices, and a volatile market makes it riskier for a company to value its shares and for investors to commit capital. Empirical evidence shows that when volatility spikes, IPO volume tends to fall. Schill (2004)[18] documents that periods of above-normal market volatility are correlated with a decline in the frequency of IPOs and the total capital raised via IPOs**.** In volatile times, companies face a higher risk that their IPO could be mispriced or even fail (if sudden market swings sour investor appetite at the last minute). Understandably, many firms elect to wait for calmer conditions rather than "sail in rough waters". Interestingly, Schill (2004) finds that higher pre-IPO market volatility did not significantly increase IPO underpricing in the deals that did go forward.

---

[17] Pástor, Ľ., & Veronesi, P. (2005). Rational IPO waves. *Journal of Finance, 60*(4), 1713–1757.
[18] Schill, M. J. (2004). Sailing in Rough Water: Market Volatility and Corporate Finance. *Journal of Corporate Finance, 10(5), 659–681.*

This suggests that issuers and underwriters are often able to adjust or delay offerings in response to volatility, rather than simply accepting much lower offering prices. In other words, when risk is elevated, the common response is to slow or pause the IPO pipeline until confidence returns, rather than proceed at any price.

The tendency to delay or cancel IPOs during bad markets is well documented in practice. Companies themselves often attribute withdrawn or shelved IPOs to unfavourable market conditions. If broad indices have been falling, or if comparable recent IPOs have fared poorly, a company planning to go public may conclude that the timing is not right. In fact, "market conditions" is the most commonly cited reason for IPO withdrawals, underscoring how integral the external environment is to the go/no-go decision. This behaviour aligns with market timing theory: firms prefer to wait rather than sell equity at a discount when the market mood is sour. Likewise, venture capitalists backing IPO-bound firms often counsel patience in volatile or bearish markets, since accepting a lower valuation or risking a failed offering could destroy value. The net result is that IPO activity slows to a trickle in bear markets or during crises, then surges when the outlook improves. For example, IPO issuance virtually stopped during the 2008–2009 financial crisis and the early 2020 COVID-induced volatility spike, only to rebound when market stability and optimism returned.

Firms not only ask "Do we need equity capital?" but also "Is now the right time to tap the market?" The evidence shows that they pay close attention to the state of equity markets, including overall returns, valuation levels, investor sentiment, and volatility in making that decision. Going public during a "hot" market wave can allow firms to raise more money on better terms, whereas attempting an IPO in a "cold" or volatile market can be costly or futile. As a result, successful corporate financing strategy often involves forecasting or gauging market conditions to choose an auspicious moment for an IPO.

1.2 Equity Capital Raising and the Decision to Go Public

When a firm seeks to raise capital for growth or other strategic needs, it must evaluate equity financing within the context of its broader financing strategy. An initial public offering (IPO), the sale of shares to public investors, is one prominent form of equity financing, but it is not the only option. Companies weigh the benefits of accessing public equity against alternatives like retaining private ownership with venture capital or private equity funding, or using debt financing. The decision to go public is thus a multifaceted strategic choice, influenced by corporate finance theory, the firm's characteristics, and the trade-offs between funding alternatives.

1.2.1 Equity Financing in Corporate Financial Strategy

Equity capital is a fundamental component of a firm's capital structure, alongside debt. From a strategic standpoint, raising equity (particularly via an IPO) can significantly strengthen a company's balance sheet by providing permanent capital that does not require repayment or incur interest. Unlike debt, equity financing carries no obligatory payments, reducing the risk of financial distress during downturns. This makes equity attractive for firms with uncertain cash flows or high growth prospects that need funding without the burden of fixed obligations. For example, young technology or R&D-intensive companies often struggle to generate steady short-term cash flows and therefore lean toward equity financing, since heavy debt would require regular interest payments that these firms might not sustain. By issuing equity, such firms obtain the necessary funds for expansion while avoiding excessive leverage. However, equity financing also comes with strategic costs. Selling stock dilutes the ownership stake of existing owners and can lead to partial loss of control or decision-making power. In a private setting, founders and early investors maintain tight control, but an IPO disperses ownership and introduces new shareholders who gain voting rights. Moreover, once public, a company must comply with extensive regulatory and disclosure requirements. The trade-off is clear: going public through an IPO provides access to deep pools of cash and liquidity, at the expense of increased transparency demands and reduced control. Remaining private (relying on private equity or debt) allows owners to retain greater operational control and confidentiality but may limit the scale of capital available. Firms must therefore assess how critical large-scale funding is to their strategy, relative to the value they place on control and flexibility.

1.2.2 Capital Structure Theories and the Go-Public Decision

Corporate finance theory offers insight into how firms choose between equity and other financing, illuminating the go-public decision. Capital structure theories such as the trade-off theory and pecking order theory are particularly relevant to equity capital raising decisions. According to the trade-off theory, firms seek an optimal debt-equity mix by balancing the benefits and costs of debt financing. Debt provides a tax shield (interest payments are tax-deductible) which can lower the firm's cost of capital, but too much debt increases the probability of costly financial distress and bankruptcy. The trade-off perspective posits that companies will issue equity (or reduce debt) when their leverage exceeds the optimal level, in order to restore balance between the tax benefits of additional debt and the rising risk of insolvency. In the context of an IPO, a firm with high leverage or constrained debt capacity might decide to go public to rebalance its capital structure by injecting equity. Indeed, evidence shows that some firms undertake IPOs not primarily to fund new investments, but to pay down debt or improve financial stability after a period of rapid growth. By raising equity capital, these firms can reduce their debt ratios and better position themselves for long-term sustainability.

Another influential framework is the pecking order theory (Myers & Majluf, 1984), which emphasizes information asymmetry between managers and outside investors. Pecking order theory states that managers prioritize funding sources in a hierarchy: first internal funds, then debt, and issue new equity only as a last resort. The logic is that managers have superior knowledge of the firm's true value and prospects, so external finance, especially equity, will be viewed with scepticism by investors who fear the firm would only sell shares if they are overpriced. As a result, issuing equity can send a negative signal. For instance, using *debt* is often interpreted as a sign that management is confident in future earnings (since debt entails fixed repayments), whereas issuing *equity* may signal that management considers the stock overvalued or the firm too risky for more debt. This adverse selection dynamic makes equity costlier and implies that firms will avoid an IPO unless other financing avenues are insufficient. In practice, many firms do follow this pattern: they tap retained earnings and private debt markets first, and turn to an IPO only when the scale of financing needed outstrips what debt can provide or when leverage has become uncomfortably high. High-growth companies, for example, often exhaust internal funds and borrowings to fuel expansion and eventually reach a point where a large equity raise via IPO becomes the only viable way to fund further growth.

While trade-off and pecking order theories sometimes predict different behaviours (the former focuses on an optimal leverage target, while the latter implies no well-defined target leverage due to financing hierarchy), both shed light on the IPO decision. A firm's *capital structure considerations* will shape its timing and willingness to go public. A company with low existing debt and strong cash flows (hence ample internal funding) might delay or avoid an IPO, consistent with pecking order preferences. Conversely, a firm facing high leverage might expedite an IPO to recapitalize with equity, aligning with trade-off theory. Empirical studies confirm elements of both theories in IPO choices. For instance, Pagano, Panetta, and Zingales (1998) find that Italian firms often went public after periods of substantial growth not to undertake new projects, but to restructure finances, reducing debt and allowing initial shareholders to diversify holdings. This suggests that beyond pure fundraising, the *decision to go public is sometimes motivated by capital structure optimization and shareholder liquidity goals*. It is also worth noting the *market timing theory* of capital raising, which posits that managers time equity issuances for windows when market conditions are favorable (e.g. high valuations or lower volatility). Firms may be more inclined to launch an IPO during bull markets or when their industry is in favour, taking advantage of potentially "overvalued" equity to minimize dilution. Such market-driven timing considerations complement the internal trade-off and pecking order factors in the go-public decision.

1.2.3 Firm Characteristics and the Likelihood of Equity Issuance

Certain firm-specific characteristics are strongly associated with the choice to raise equity and pursue an IPO. Size and age are important determinants: IPO candidates are often of a sufficient size and at a stage of development where they require capital beyond what private sources can supply. Young firms tend to have higher growth opportunities and capital needs, which push them toward equity issuance, but they must also reach a credible size and track record to attract public investors. Very early-stage startups usually rely on private financing until they achieve the scale for a successful IPO. Empirical research shows that firm age and size influence the timing of going public consistent with a life-cycle view that companies go public after achieving a certain maturity and when expected benefits outweigh the costs. At the same time, being too mature or asset-rich can make debt financing easier (due to collateral and steady earnings), so mid-stage high-growth firms are often the prime IPO candidates.

Growth prospects and intangible assets also drive the preference for equity. Companies with strong growth opportunities (e.g. in technology or biotech sectors) often need substantial funds for expansion, research, and development. Their intangible assets (such as intellectual property or human capital) cannot be readily used as collateral for loans, limiting their debt capacity. High-growth, R&D-intensive firms thus favour equity financing because it does not constrain them with fixed payments and scales with their valuation potential. As noted, these firms might otherwise face a debt funding gap due to lack of collateral and the riskiness of their projects, making an equity infusion via IPO an attractive solution. In contrast, firms with more stable cash flows or tangible assets (e.g. in manufacturing or utilities) can often sustain higher debt, and may have less frequent need for public equity infusions.

The presence of *venture capital (VC) or private equity backing* is another key characteristic linked to IPO decisions. Venture-backed firms frequently use IPOs as an exit strategy for early investors and a stepping stone for further growth. Venture capitalists typically invest with the expectation of eventually liquidating their stake, and an IPO provides a liquid market for the shares. Moreover, VC-backed companies often benefit from the credibility and professionalization that VCs bring, making them well-prepared for public markets. Studies indicate that VC involvement can certify firm quality and reduce information asymmetry at IPO, which may lower underwriting costs or underpricing. From a strategic standpoint, a firm backed by venture capital is more likely to go public once it achieves a strong market position, as the IPO allows VCs to realize returns and the company to access a broader investor base. Similarly, companies backed by private equity (PE) sponsors (in later-stage buyouts) may pursue IPOs to partially exit and deleverage. In Europe, many IPOs, especially larger ones, are "sponsor-backed" offerings where a PE firm sells down its stake. Such firms are often characterized by already sizeable operations and profitability, but use the IPO to achieve a lower debt load and diversify ownership.

Other characteristics include *financial performance and outlook*. Firms with high past growth and optimistic future earnings forecasts are more inclined to choose equity issuance, as public investors are willing to pay for that growth potential. Profitability, on the other hand, can cut both ways: highly profitable firms generate internal cash (reducing immediate need for equity per the pecking order theory), but profitability also reassures public investors, making an IPO more feasible. Thus, moderately profitable, high-growth firms may be ideal IPO candidates; they have enough track record to appeal to the market but still have ample investment needs. Finally, ownership structure can influence the go-public choice. Founder-led or family firms might be reluctant to dilute control through an IPO unless necessary, preferring family funds or bank loans. By contrast, firms with more diffuse private ownership or those needing to distribute wealth (e.g. multiple founders or early employees looking to liquidate some shares) will find an IPO attractive to provide liquidity for existing shareholders. In summary, companies that are younger (but established), growing quickly, asset-light, and investor-backed are generally more predisposed to raise equity publicly, whereas older, asset-heavy firms or those with easy access to credit may delay or avoid going public.

1.2.4 Strategic Motivations for Going Public Beyond Capital Raising

In evaluating an IPO, firms consider not only the immediate influx of capital but also a range of *strategic benefits* that come with being publicly traded. While raising capital for growth or debt reduction is the primary objective, several *ancillary motivations* often factor into the decision:

- *Creating a currency for acquisitions*: An IPO enables a firm to have publicly traded stock that can be used as "currency" to finance mergers and acquisitions. CFO surveys show that the ability to use stock for acquisitions is frequently ranked as one of the most important reasons to go public. Acquisitive companies value a public listing because it allows them to more easily buy other firms by paying in stock rather than cash. Smaller firms, in particular, gain a powerful tool post-IPO to pursue strategic acquisitions for growth.
- *Liquidity and diversification for owners*: Going public creates a liquid market for shares, addressing the inherent illiquidity of private company stock. Founders, early investors, and employees can sell portions of their holdings over time, achieving personal diversification and exit opportunities. This liquidity is a strong motivator for venture capitalists and founders who have much of their wealth tied up in the company – an IPO allows them to realize gains and reduce risk, even if they retain a significant stake. In addition, the public market's liquidity means the company can more readily raise follow-on equity in the future if needed, since investors are more willing to buy shares that they know can be resold.
- *Enhanced corporate visibility and credibility*: A public listing often brings prestige and can heighten a company's profile with customers, business partners, and the media. The stringent disclosure and governance requirements for public firms signal transparency and stability, which can make counterparties more willing to do business. Being publicly traded thus can serve as a signalling mechanism of the company's quality and legitimacy. For example, the IPO process itself subjects a firm to scrutiny by regulators and investors, and a successful offering suggests that the firm has passed those market tests. This credibility can translate into tangible benefits like better credit terms from banks, stronger negotiating power with suppliers, or increased ability to attract top talent.
- *Improved corporate governance*: Transitioning to public ownership typically involves strengthening governance structures; independent board members, formalized reporting systems, and adherence to market regulations. These changes can help mitigate agency problems by imposing more discipline on management. Shareholders of public companies can monitor performance through stock price movements and have a voice via shareholder meetings. Some firms pursue IPOs in part to institutionalize governance practices and impose market discipline on management for the long run. This can lower the cost of capital, as investors perceive lower agency risk in a firm with transparent operations and oversight.
- *Employee incentives and hiring*: Public companies can more readily implement stock-based compensation plans (such as stock options or restricted stock units) that attract, retain, and motivate employees. In fast-growing sectors, talented employees often seek the upside of equity participation. An IPO makes the company's equity a liquid, market-valued currency, which can be a powerful incentive. It enables employees to share in the company's success and aligns their interests with shareholders. Moreover, having a visible stock price can foster a sense of accountability and performance orientation among employees.

These strategic considerations mean that the decision to go public is not solely about raising immediate funds. It is also about leveraging the public status to support the company's long-term strategy; whether that is growth through acquisitions, providing exits for investors, enhancing reputation, or improving governance and operations.

1.2.5 Funding Alternatives and Trade-Offs

Before committing to an IPO, companies typically explore alternative financing paths, each with its own trade-offs. *Private equity financing* (including venture capital for earlier-stage companies and growth equity or buyout funds for more mature firms) is a common route. Private investors can inject substantial capital without the firm facing public-market scrutiny. This route allows the company to remain private longer, often under the guidance of experienced investors who may help professionalize the business. The trade-off is that private equity investors demand significant ownership stakes and influence, and often expect a higher return on investment (reflected in the eventual need for an exit like an IPO or sale). Taking on a private investment can delay the IPO, but it does not eliminate the eventual pressure for liquidity; in many cases it actually sets the stage for a larger IPO down the line. Firms must consider *whether the loss of ownership and control to a private investor is preferable to the wider ownership but regulatory burdens of an IPO.*

Increasing *debt financing* is another alternative to raising equity. Companies with strong cash flows or collateral might borrow from banks or issue corporate bonds to meet funding needs instead of diluting equity. Debt can be cheaper in terms of cost of capital (due to tax deductibility of interest) and keeps ownership intact, but it also heightens financial risk. High leverage constrains a firm's flexibility and can lead to distress if earnings falter. Thus, firms with volatile or uncertain prospects are limited in how much debt they can prudently take on. Additionally, debt often comes with covenants and restrictions that could hinder strategic moves. The IPO decision may hinge on these considerations: if taking more debt would overstretch the balance sheet or impose too many restrictions, the firm may opt for equity despite the dilution. In essence, companies evaluate *debt vs. equity by comparing the marginal cost of borrowing against the opportunity cost of issuing shares*, when debt's risks outweigh its benefits, equity becomes the favoured option, and vice versa.

Some firms also weigh *strategic partnerships or minority stakes* from larger corporations as a financing avenue. A strategic investor might provide capital in exchange for equity and possibly some form of alliance or collaboration. This can offer not only funding but also business synergies or market access. The downside is that it might limit a company's strategic flexibility (if tied to one partner) and the amount of capital is often less than what an IPO could raise. Furthermore, a significant corporate investor on the cap table could deter other investors or acquirers. Compared to an open-market IPO, private strategic investments are typically negotiated and may undervalue the company's upside in exchange for the strategic benefits offered.

Finally, a major alternative to going public is *outright acquisition*: rather than raising capital independently, a company might choose to be acquired by a larger firm with substantial financial resources. This provides liquidity to owners and can fund growth as the acquirer takes over financing future investments. The trade-off here is the loss of the company's independence and the founders' vision being subsumed under the parent company's strategy. Many entrepreneurs consider an IPO preferable if they wish to continue leading an independent company, whereas selling to a strategic acquirer or a private equity buyer means ceding control. Thus, the choice between an IPO and a sale or private funding often boils down to the founders' objectives and the company's long-term strategic plan.

In weighing these alternatives, firms perform a *cost-benefit analysis under uncertainty.* They consider the relative cost of capital (e.g. interest rates vs. equity cost), the implications for control and governance, and the constraints each option imposes. The optimal path can also be contingent on market conditions, for example, if stock market valuations are very high, an IPO may yield more capital for less dilution (tilting the trade-off in favour of going public). Conversely, in times of market volatility or recession, even companies inclined to go public might postpone and use interim private funding to bridge the period. Thus, the *go-public decision is dynamic*, responding to both internal needs and external market climates.

1.2.6 Capital Structure Considerations in the Go-Public Decision

Ultimately, a firm's decision to raise equity capital through an IPO is shaped by how it *balances capital structure considerations with growth ambitions and risk tolerance.* The firm's current leverage, financing needs, and the theories discussed all inform this balance. If a company's capital structure is overly skewed to debt, the trade-off theory would encourage issuing equity to prevent financial distress, making an IPO an attractive solution if private equity is insufficient or undesirable. If a company can comfortably fund projects through retained earnings and debt (and is wary of adverse signals), pecking order theory implies it will delay equity issuance, perhaps until *a strategic inflection point* is reached where equity is unavoidable (such as a transformative expansion or when debt capacity runs out). Firms also consider their *target* capital structure post-IPO, many companies explicitly use IPO proceeds to pay down debt and achieve a more conservative leverage ratio ahead of being public, to ensure financial flexibility and appeal to a broader investor base.

In practice, going public is often the culmination of a financing journey where internal funds and private financing have been utilized to the extent reasonable. The IPO then injects a large pool of capital that can dramatically deleverage the firm or fund big investments, resetting the capital structure for a new stage of growth. Notably, once public, the company has ongoing access to equity markets (for secondary offerings) and often improved access to debt as well, due to the transparency and market valuation in place. This means the initial decision to go public is also a decision to *embrace a new financing paradigm*: one where the firm's capital structure can be actively managed with a mix of public debt and equity, rather than being constrained to private sources.

Before proceeding, managers must be cognizant of the risks and costs involved. Despite the many benefits of IPOs outlined, they come hand-in-hand with significant risks that can influence a firm's willingness to go public, especially under volatile or uncertain conditions. For instance, once public, the company's valuation will be subject to market fluctuations beyond management's control. *Exposure to market conditions* means that even a fundamentally strong firm could see its stock price decline due to external shocks or sector-wide downturns, potentially making it a target for hostile takeovers or causing pressure from activist shareholders. The IPO process itself carries the risk of underpricing (selling shares below their true value), substantial transaction fees, and possible failure if market sentiment sours at the wrong moment. These considerations often give firms pause, leading them to carefully time their IPOs or include protective strategies (such as staggered share lock-ups or conservative pricing).

In summary, the decision to utilize equity capital through an IPO is a complex strategic choice. Firms must integrate capital structure theory, their own characteristics, and strategic objectives to determine if and when going public maximizes shareholder value. While an IPO can unlock growth opportunities, diversify ownership, and improve financial flexibility, it also exposes the firm to the disciplines and whims of the public market.

1.3 Risks Faced by Firms During an IPO

In Section 1.2, we discussed how firms strategize their financing and why they may choose to raise equity through an Initial Public Offering (IPO). While an IPO can provide substantial capital and other strategic benefits, it is a complex process fraught with multiple risks. Volatile market conditions in particular can amplify these risks, affecting both the process and outcomes of going public. Firms pursuing public equity financing face three key forms of risk exposure: offer price risk**,** distribution risk, and aftermarket risk**.** These correspond to uncertainty in (i) the IPO's offer price, (ii) the ability to sell the offered shares to investors, and (iii) the stock's performance and liquidity after listing, respectively. Firms considering an IPO will only proceed when they expect that the long-term benefits outweigh these risks – a calculus that becomes especially challenging in volatile market environments. This section unpacks each of these risks in turn, overall, the goal is to demonstrate that IPOs, while potentially valuable financing avenues, carry risks that interact with market conditions (especially volatility) in important ways.

1.3.1 Offer price risk

Offer price risk is the uncertainty surrounding the price at which the firm's shares can be sold in the IPO. Unlike a seasoned offering where a current market price exists, an IPO involves discovering a new market price for the company's equity. Setting the "right" offer price is challenging due to information gaps and market fluctuations. In practice, the IPO process involves substantial lead time (often several months) for preparatory steps such as due diligence, regulatory filings, and investor marketing. During this time, market conditions can change significantly, causing the valuation of the firm to rise or fall. If market prices increase in the interim, the issuer may be able to raise the offer price and reap greater proceeds; conversely, if market prices decline, the firm might be forced to cut the IPO price or even postpone/cancel the offering. This exposure to market movements before the pricing date constitutes offer price risk. Both the issuing firm and its underwriters bear this risk, since underwriters' fees are usually proportional to the offering price. A sharp adverse price move can therefore reduce the funds raised and the underwriter's compensation, or render the IPO unattractive to proceed. Indeed, in volatile markets, firms commonly withdraw or delay IPOs if the expected price falls below acceptable levels. Regulatory provisions (such as Rule 477 in the U.S.) allow issuers to withdraw a registered offering, and exercising this option, while costly in terms of wasted effort and underwriting expenses, may be preferable to proceeding under unfavourable pricing. Empirical evidence confirms that issuers cite "unfavourable market conditions" or a low expected price as the primary reason for withdrawing IPOs. For example, during a spike in market volatility in late 1998, 18 out of 20 U.S. firms that withdrew their IPO filings explicitly blamed adverse market conditions. This underscores how volatile equity markets magnify offer price risk and can derail equity financing plans.

A central driver of offer price risk is *information asymmetry* between the firm and investors regarding the company's true value. In IPOs, some investors (typically institutional or informed investors) have better information about the firm's prospects, while others (the general public or uninformed investors) do not. This information imbalance gives rise to the classic *winner's curse* problem in IPO pricing. Rock (1986)[19] formalized the winner's curse model, showing that informed investors will bid for shares only when an IPO is underpriced (i.e. when they perceive the offer price is below the stock's intrinsic value), whereas uninformed investors risk ending up with the bulk of shares in overpriced or low-quality offerings. As a result, issuers are compelled to underprice IPOs on average to attract sufficient demand from uninformed investors. In other words, underpricing serves as compensation for the winner's curse adverse selection, ensuring that IPO participation remains attractive for all investor groups. Consistent with this theory, Beatty and Ritter (1986)[20] find that IPO underpricing tends to be higher for firms with greater ex ante uncertainty (risk), implying that when the true value of the company is harder to pin down, issuers leave a larger "money on the table" to entice investors, effectively a risk premium for the offer price uncertainty. Information asymmetry-driven pricing risk is also reflected in the *partial adjustment phenomenon*: when strong investor demand is revealed (e.g. through book-building), issuers revise prices upward, but typically not by the full amount of the positive information, resulting in a high first-day return; this suggests issuers and underwriters are cautious in fully adjusting the offer price, partly to insure against estimation errors and to reward investors for information (Hanley, 1993) [21].

Related to information asymmetry is the concept of *signalling* in IPO pricing. High-quality firms may deliberately underprice their IPO as a signal of quality to the market. The intuition is that a company confident in its future prospects can afford to sell shares at a discount today, knowing it will be able to issue additional equity later at higher prices, whereas lower-quality firms cannot easily mimic this strategy. Several theoretical models formalize this signalling argument: for

---

[19] Rock, K. (1986). Why new issues are underpriced, *Journal of Financial Economics,* 15, pp. 187–212.

[20] Beatty, R. P., & Ritter, J. R. (1986). Investment banking, reputation, and the underpricing of initial public offerings. *Journal of Financial Economics*.

[21] Hanley, K. (1993). The underpricing of initial public offerings and the partial adjustment phenomenon. *Journal of Financial Economics, 34*, 213–250.

instance, Allen and Faulhaber (1989)[22], Grinblatt and Hwang (1989)[23], and Welch (1989) [24] show that underpricing can serve as a credible signal of firm quality. Empirically, some studies have found patterns consistent with signalling (e.g. firms that heavily underprice might subsequently return to the market with successful seasoned offerings) though the evidence is mixed. It is worth noting that underpricing as a signal or as winner's curse compensation is a cost to the original owners (since they receive less proceeds). However, they may rationalize it as an investment in the firm's market reputation or as insurance against a failed issue. In fact, Ibbotson (1975)[25] had earlier suggested that issuers underprice IPOs to "leave a good taste in investors' mouths" hoping to cultivate goodwill that benefits future capital-raising. This reputational notion overlaps with signalling and highlights that offer price risk is not just about setting the single right price, but *managing the market's reception* of the IPO. Other early hypotheses posited that underpricing could be a form of *insurance against legal liability*, by ensuring the stock trades up, the issuer and underwriters reduce the likelihood of investor lawsuits alleging misrepresentation. Tinic (1988)[26] models this "insurance" motive and indeed argues that underpricing lowers the probability of litigation, especially in markets with high legal liability risk. Overall, the literature indicates a variety of reasons why IPO offer prices might deviate from the issuer's true reservation value. What unites these perspectives is the inherent uncertainty in estimating market demand and value for a new issuance. Offer price risk is thus a fundamental challenge: *if the price is set too high, the issue may flop (or require withdrawal); if set too low, the firm leaves money on the table and potentially signals information (rightly or wrongly) about its quality.*

It is important to connect offer price risk with broader market volatility. High market volatility exacerbates offer price risk by making the pricing outcome more uncertain. When overall stock prices swing widely, the value of the firm (and comparables) can change materially between the decision to go public and the actual pricing. The *option to withdraw or delay* the IPO provides a partial hedge against this risk, effectively allowing issuers to wait for more favorable conditions rather than locking in a depressed price. Busaba et al. (2001)[27] find that the *threat* of withdrawal during book-building can even strengthen the issuer's bargaining position and lead to less underpricing, since investors know the issuer will walk away if the indicated price is too low. Thus, issuers can manage offer price risk to an extent, but not eliminate it. Empirical patterns in IPO

---

[22] Allen, F., & Faulhaber, G. R. (1989). Signalling by underpricing in the IPO market. *Journal of Financial Economics, 23*(2), 303–323.
[23] Grinblatt, M., & Hwang, C. Y. (1989). Signalling and the pricing of new issues. *Journal of Finance, 44*, 393–420.
[24] Welch, I. (1989). Seasoned offerings, imitation cost, and the underpricing of initial public offerings. *Journal of Finance, 44*, 421–448.
[25] Ibbotson, R. (1975). Price performance of common stock new issues. *Journal of Financial Economics, 2*, 235–272.
[26] Tinic, S. M. (1988). Anatomy of initial public offerings of common stock. *Journal of Finance, 43*, 789–822.
[27] Busaba, W. Y., Benveniste, L. M., & Guo, R. (2001). The option to withdraw IPOs during the premarket: Empirical analysis. *Journal of Financial Economics, 60*, 73–102.

timing support the view that firms attempt to time their offerings for periods of lower risk: for example, studies show IPO volume tends to decline in times of high volatility and to resume when volatility abates or markets have risen. Schill (2004) reports that periods of abnormally high market volatility are associated with a significant drop in the number of IPOs and amount of proceeds raised, as firms become wary of the price risk in turbulent markets. In Europe as well, IPO activity has been observed to wax and wane inversely with volatility, and many European issuers pulled offerings amid the uncertainty of the 2008 financial crisis and other volatile episodes. In summary, offer price risk encompasses the *uncertainty in discovering the proper IPO price*, driven by information asymmetry and market fluctuations. It is typically manifested in IPO underpricing (a deliberate outcome to mitigate risk) or in withdrawn/failed offerings when the pricing cannot be resolved favourably. This risk is a first-order consideration for any firm contemplating an IPO, as it directly impacts the firm's financing outcome and cost of capital.

1.3.2 Distribution risk

The second major challenge is distribution risk, the risk that the IPO may not attract enough investor demand to sell the full offering at the intended price. Even after an offer price is set, the issuer faces uncertainty about placing the shares with investors. If the offering is undersubscribed (demand falls short), the issuer may have to cancel the issue or allow the price to drop to entice buyers. Distribution risk is closely related to offer price risk, but it emphasizes the *quantity of shares sold (y)* rather than the price per share. Schill (2004) defines distribution risk as the uncertainty associated with the number of shares that can be sold at the offer price. In practice, the distribution phase of an IPO, typically the book-building period and the actual allocation of shares, can last from days to a couple of weeks. During this time, market conditions and investor sentiment play a critical role. If markets turn volatile or negative, investors may pull back orders, *increasing the probability that the issue cannot be fully sold at the set price*. Schill notes that the probability of not selling the entire offering at the offer price *becomes larger during periods of greater price volatility*. In other words, volatile markets heighten distribution risk: even if an initial price seems fair, swift changes in sentiment can leave an IPO under-subscribed.

The book-building process used in most modern IPOs (including virtually all European IPOs in recent decades) is designed to manage distribution risk by gauging investor interest in advance. During book-building, underwriters solicit indications of interest from institutional investors, effectively building a demand curve for the shares. This allows the underwriter to adjust the price or the size of the offering to match demand, reducing the risk of a failed distribution. Nonetheless, book-building does not eliminate distribution risk; it merely provides information and flexibility. If demand is weak, the underwriter might revise down the price or encourage the firm to scale back the offering. In extreme cases, the issuer may withdraw the IPO if book-building reveals insufficient demand at any reasonable price (a scenario often foreshadowed by tepid investor feedback). *Underwriter incentives and contractual arrangements* are central to distribution risk. In a firm-commitment underwriting (the most common arrangement in the U.S. and Europe), the underwriter purchases the entire issue from the company and then resells it to investors, this means the underwriter bears the distribution risk, as the issuer is guaranteed the proceeds (minus fees) if the IPO goes ahead. Under firm commitment, the underwriter has a strong incentive to ensure all shares can indeed be sold (or else the underwriter is left holding unsold stock). To manage this, underwriters often form syndicates (groups of banks/brokers) to widen the distribution network and spread the risk. Empirical evidence shows that for larger offerings, which are harder to place, underwriters use larger syndicates, implying that distribution risk increases with the scope of the issue. For instance, a study of Brazilian IPOs found that syndicate membership grew with offer size, "suggesting that distribution risk was relevant" in necessitating more underwriters (Almeida & Leal, 2015)[28].

By contrast, in a *best-efforts* contract, the underwriter does not guarantee the sale of all shares but only agrees to use their best effort to sell what they can; here the issuing firm itself bears the distribution risk (unsold shares mean the firm raises less money). Best-efforts IPOs are relatively rare for mainstream offerings, precisely because they leave the issuer vulnerable to financing shortfalls. Historically, best-efforts contracts have been used for smaller or riskier offerings (such as some small-cap or penny-stock IPOs), where underwriters are unwilling to take on full risk.

Distribution risk also helps explain IPO underpricing from another angle. Underpricing can be viewed as a tool to ensure that the issue is easily distributed: by pricing shares at a discount to expected market value, underwriters create excess demand, virtually guaranteeing a fully subscribed issue and a successful sale. Baron (1982)[29] presented a model where issuers, recognizing their own informational disadvantage relative to underwriters, allow underwriters discretion to set a lower offer price; the underwriter in turn uses the underpricing to more easily sell the shares and reduce distribution effort. In this sense, underpricing is partly a *buffer against distribution failure*. When an IPO pops (surges) on the first day due to underpricing, it indicates demand exceeded supply – a deliberate outcome that ensures all the shares could be placed with eager buyers. If instead the IPO price is set too aggressively (no discount), there is a risk that demand will not fully materialize, leaving the underwriter and issuer in a difficult position. Underwriters strongly prefer to avoid "cold" IPOs (where the stock falls below the offer price), as those reflect insufficient demand and can damage the reputations of both the underwriter and the issuer. In fact, *underwriters often engage in price stabilization in the immediate aftermarket* to mitigate distribution shortfalls: they may buy back shares or use the *overallotment (greenshoe) option* to prop up a weakly received issue. Typically, underwriters intentionally oversell (allocate more shares than originally issued, up to 15% more) and begin the first day of trading with a short position in the stock. If the IPO is trading poorly (below the offer price), the underwriter covers this short by buying shares in the market, thereby supporting the price and ensuring that any remaining unsold allotment is purchased. This is effectively an *inventory risk management* technique: the overallotment option gives underwriters leeway to absorb excess supply or inject demand to prevent a price collapse. Such practices show

---

[28] de Souza e Almeida, V., Leal, R.P.C. (2015). Brazilian initial public offerings, underwriters, and premium corporate governance segments listing. *Corporate Ownership & Control*, 13(1-11), 1410-1418

[29] Baron, D. P. (1982). A model of the demand for investment bank advising and distribution services for new issues. *Journal of Finance, 37*, 955–976.

that the boundary between distribution risk and aftermarket risk is blurry, *actions taken to ensure distribution (like underpricing and stabilization) also influence the aftermarket performance***.**

It is worth noting that distribution risk is conditioned by *investor sentiment and market trends*. During "hot issue" markets (bullish periods with exuberant sentiment), distribution risk is lower, many IPOs get oversubscribed, and issuers can afford to raise prices. In contrast, during "cold" markets or high volatility, even decent companies might struggle to find buyers, raising the risk of undersubscription. For example, in the aftermath of financial crises or during sudden volatility spikes, IPO book-building often reveals a lack of appetite, leading to postponements. Chambers and Dimson (2009)[30] document that IPO underwriting in Europe has evolved to include stronger *commitment provisions* from underwriters precisely to reassure issuers on distribution success, and techniques like anchor investments (cornerstone investors) are sometimes used to secure a base level of demand before an IPO. All these measures indicate that the *core concern of distribution risk is to ensure that the offering can be completed in full*. A failed or significantly downsized IPO can have serious consequences: the firm may not receive the needed capital and could suffer a credibility loss in the market. Therefore, both issuers and underwriters are motivated to mitigate distribution risk through pricing strategy, marketing efforts (roadshows to generate interest), and contractual arrangements.

1.3.3 Aftermarket risk

Even after a successful issuance (i.e. the IPO is priced and all shares are sold), a firm must contend with aftermarket risk, which refers to the uncertainty in the stock's price performance, volatility, and liquidity after the IPO. In other words, once the company is public, how will its shares trade in the open market, and what risks does that pose to the firm? A newly listed stock often experiences *high volatility in early trading* as the market digests new information and as initial allocation imbalances get sorted out. Unlike established companies, an IPO stock has no prior trading history, and its investor base might be still forming; this can lead to significant price swings in the days and months following the listing. Aftermarket risk is a concern for issuers because extreme volatility or a sharp price drop post-IPO can damage the company's reputation and its ability to raise capital in the future. It may also reflect underlying *illiquidity*, if the IPO float (the proportion of shares publicly traded) is small, the stock's trading can be thin, resulting in large price impacts from relatively small trades. Illiquidity and volatility tend to go hand-in-hand, and smaller IPOs or those with concentrated ownership often face this risk until a broader market develops for their stock.

One aspect of aftermarket risk is the possibility of *underperformance* or declines in the stock price after the initial euphoria. Numerous studies (mostly focused on the U.S. market) have found that on

---

[30] Chambers, D., & Dimson, E. (2009). IPO underpricing over the very long run. *Journal of Finance, 64*(3), 1407–1443.

average, IPO stocks underperform comparable firms or market indices in the long run (three to five years after listing) – see, for example, Ritter (1991) and Loughran & Ritter (1995) for evidence of this phenomenon. Several European studies have reported similar patterns of long-run underperformance for IPOs, though results vary by country and era. For the issuing firm, a collapsing share price in the aftermarket is problematic: it can imply that the IPO was overpriced or that the firm's prospects have soured, and it may close off avenues for raising additional equity capital (a secondary offering) until the stock recovers. In extreme cases, a post-IPO price collapse can invite class-action lawsuits (in litigious markets like the U.S.) or activist pressure on management. Thus, firms are cognizant of aftermarket volatility as a risk, they do not simply maximize the IPO price without regard to what comes next. The *IPO pricing and allocation strategy often anticipates aftermarket trading*: for instance, underwriters may favour a mix of long-term institutional investors in the allocation to support the stock, or include lock-up agreements preventing insiders from selling immediately, so as to avoid a flood of supply in the early aftermarket that could depress the price.

Aftermarket risk also encompasses reputational concerns for both issuers and underwriters. Ibbotson's (1975) "reputational capital" hypothesis and subsequent work suggest that companies care about post-IPO performance because it signals their *quality and credibility* to the market. A strong aftermarket performance (the stock trading above the issue price and showing stability) leaves investors with a "good taste" potentially making them more receptive to future offerings by the company (e.g. a follow-on equity issue or even just more goodwill toward the firm's products and services). Conversely, if an IPO breaks issue (trades below the offer price) or experiences high volatility, investors may become wary of the firm, harming its financing reputation. Underpricing is one tool issuers use to tilt the odds of a favorable aftermarket: by leaving a margin for a first-day run-up, they improve the chances that the stock will trade above the offer price in early days, which is often perceived as a successful IPO. This relates back to the earlier discussion of signalling and leaving a good taste, it is essentially about managing aftermarket perceptions. Schill (2004) notes that building a financing reputation via strong aftermarket performance becomes more challenging in volatile markets, as unpredictability makes it harder to ensure a stable upward trajectory. During periods of high market volatility, even a fundamentally solid company's stock might whipsaw due to external factors, undermining the firm's attempt to signal quality through a smoothly rising post-IPO price. Thus, volatility adds a layer of hazard to reputational strategy in the IPO context.

Underwriters, too, have reputational stakes in the aftermarket. Prestigious investment banks are often associated with "quality" IPOs, and they vigilantly guard against their IPOs performing poorly once trading begins. If a top-tier underwriter's deals frequently collapse in the aftermarket, that underwriter will have trouble attracting future IPO mandates (issuers would question their pricing and placement skill). Investors also recall which banks bring IPOs that crash, those underwriters might find it harder to place the next deal. In response, underwriters engage in *price support* (as discussed earlier) and are selective about the deals they take public during volatile times. Interestingly, research has found that reputable underwriters actually scale back their IPO activity when market volatility is high (suggesting they avoid exposing their reputation to high aftermarket risk). Wolfe, Cooperman, and Ferris (1994)[31] document that prestigious underwriters are less active in volatile markets, consistent with the idea that they prefer to wait out "rough waters" rather than launch IPOs that might fail and tarnish their record. All of this underscores that aftermarket risk is not an afterthought, but rather anticipated and managed by market participants as part of the IPO process.

Key components of aftermarket risk include short-term volatility, trading volume (liquidity), and longer-term stock performance. Firms and underwriters aim for an orderly market debut, for example, using stabilization measures to prevent a free-fall in the days after listing. In many European IPOs, especially larger ones, there is an overallotment (greenshoe) facility and sometimes formal stabilization agents appointed for a few weeks post-IPO to reduce volatility. Despite such efforts, the reality is that once the stock is on the exchange, its price is subject to market forces beyond the issuer's control. Broader market volatility will affect the stock; if the IPO coincides with turbulent conditions, its aftermarket path can be rockier. Recent studies highlight that market volatility around the time of an IPO can adversely impact post-IPO liquidity and increase price volatility, thereby raising the cost of capital for the firm (since investors demand a higher return for holding a riskier, more volatile stock). Moreover, the *post-IPO performance has implications for investor confidence*: for instance, headlines about weak post-IPO performance or high volatility can sour sentiment and disincentivize other companies from going public in the same window. Thus, a volatile aftermarket can not only hurt the firm in question but also reflect wider conditions that feed back into IPO market cycles (hot vs. cold issue periods).

In summary, *aftermarket risk* encompasses the uncertainties a firm faces once it has become publicly traded, will the stock price be stable or volatile? Will it sustain or exceed the IPO valuation, or will it drop and disappoint investors? Will there be sufficient liquidity for investors to trade without large price impacts? These questions matter to firms because they influence the success of the IPO in a broader sense than just closing the deal, they affect the firm's ongoing access to capital and public market reputation. Companies that navigate the IPO only to see their stock flounder may find that the supposed benefits of going public (better access to capital, heightened market profile, etc.) are slow to materialize or even reversed. Therefore, prudent issuers (and their underwriters) plan for the aftermarket: they might, for example, ensure that a diverse and quality set of investors are in the book (to avoid one-sided order flow later), or they communicate a steady news flow post-IPO to maintain interest in the stock. The interplay between aftermarket risk and prior stages is clear, choices made in pricing and allocation (offer price and distribution) directly impact the aftermarket. For instance, an overvalued, just-barely-placed IPO is likely to struggle afterward (as savvy investors sell off, driving the price down), whereas a moderately underpriced, highly sought-after IPO tends to enjoy a positive after-market momentum.

[31] Wolfe, G. A., Cooperman, E. S., & Ferris, S. P. (1994). An analysis of the underwriter selection process for initial public offerings. *Journal of Financial Research, 17*(1), 77–90.

1.4 The Influence of Market Conditions on Equity Issuance

In the previous section, the various risks faced by firms during an IPO were examined, highlighting how uncertainty can complicate the going-public process. A key external source of such uncertainty is the broader market environment in which an IPO takes place. Market conditions, encompassing stock market volatility, overall equity price levels, investor sentiment, interest rates, and macroeconomic trends, critically shape a firm's decision to issue equity and the ultimate success of that issuance. Firms do not make IPO timing and pricing choices in isolation; rather, they respond strategically to the prevailing capital market climate, seeking favorable "windows" while guarding against adverse conditions. This section explores how these market factors influence IPO volume, timing, pricing, and outcomes, blending theoretical models with empirical evidence to illustrate the interaction between firm decision-making and broader capital market dynamics.

1.4.1 Theoretical Perspectives: Asymmetric Information and Market Timing

Classic finance theory suggests that information asymmetry and market timing considerations drive equity issuance decisions. In asymmetric information, managers are hesitant to issue new shares if they believe the firm is undervalued, because issuing under unfavourable conditions dilutes value for existing owners. In contrast, they are more inclined to issue equity when valuations are high or possibly overvalued, to capitalize on favorable pricing. Consistent with this logic, evidence indicates that companies tend to successfully time their offerings for periods when valuations are high meaning firms purposely go public during bullish markets when they can fetch better prices. Such market timing behaviour implies that IPO waves often coincide with periods of strong stock market performance and optimistic valuations. In fact, one influential model finds that *IPO waves* are typically *preceded by high market returns and followed by lower returns*, suggesting firms rush to issue equity after stock prices have risen and the cost of capital has fallen; this timing may be perfectly rational, a decline in expected market returns or an increase in expected profitability can trigger an optimal wave of IPOs, but it also opens the door for less fundamentally justified issuances if investors are overly optimistic (Pastor and Veronesi, 2005).

Investor sentiment and behavioural factors offer an alternative (though complementary) perspective on market timing. During "hot issue" markets, investor exuberance can reduce the effective cost of equity by boosting demand for new offerings, thereby encouraging more firms to go public. High sentiment periods often feature both *high IPO volume* and *high initial returns* (underpricing), as enthusiastic investors rush to buy into new issues. However, theory and evidence suggest that these hot-market issuers may not always be the highest-quality firms. Asymmetric information models predict that in good times even riskier firms can issue because the perceived "lemons" discount is lower. Empirical research by Helwege and Liang (2004)[32] supports the view that hot IPO markets are driven in part by investor overoptimism: firms going public in booming markets exhibit

[32] Helwege, J., & Liang, N. (2004). Initial public offerings in hot and cold markets. *Journal of Financial and Quantitative Analysis, 39*(3), 541–569.

similar ex-ante operating performance to those in cold markets, yet their stocks underperform subsequently, which is largely consistent with investor overoptimism in hot markets, but not with the asymmetric information models. In other words, buoyant market conditions can induce investors to overvalue new issues, and managers respond by taking advantage of such sentiment. This behavioural dimension reinforces why market conditions, beyond just fundamentals, have a powerful influence on equity issuance decisions.

1.4.2 Empirical Patterns: Volatility, Sentiment, and Macroeconomic Climate

A substantial body of empirical evidence links financial market variables to IPO activity. Perhaps the most direct factor is market volatility. Volatility is a proxy for uncertainty and risk in the market; high volatility tends to unsettle investors and issuers alike. Research shows a clear negative relationship between volatility and IPO issuance. In volatile environments, even firms that proceed with an IPO must contend with greater pricing risk and possibly higher issuance costs (underwriters charge higher fees to compensate for distribution risk). Notably, increased volatility does not necessarily increase IPO underpricing, suggesting that issuers and underwriters may delay offerings rather than massively discount the price when markets are turbulent. This evidence aligns with Mandelker and Raviv's (1977)[33] risk-bearing model, where underwriters and issuers are reluctant to float new equity in "rough waters" unless compensated. In short, heightened volatility dampens IPO activity, particularly for smaller or less seasoned firms that are most vulnerable to market swings.

Beyond volatility, the *level of stock market indices and recent market returns* are critical barometers for IPO timing. IPO markets are often described as cyclical, closely tied to broad equity performance and investor appetite. When major stock indices are rising and valuations are rich, companies find the window of opportunity to issue equity much more inviting. Empirical analyses have identified market sentiment and valuation levels as key determinants of IPO volume. Lowry (2003)[34], for instance, shows that the two most important factors behind fluctuations in IPO volume are firms' demand for capital (often higher during economic expansions) and investor optimism. Periods of robust economic growth and bull markets often coincide with increased capital needs (firms seek funds for new projects) *and* a receptive investing public, a recipe for an IPO boom. In contrast, if stock prices are depressed or investors are pessimistic, even firms in need of capital may postpone offerings due to the high adverse-selection costs of issuing equity under those conditions. Lowry finds that adverse-selection costs (information asymmetry) do play a statistically significant role in IPO fluctuations, but their economic impact is relatively small compared to broader market sentiment. This suggests that while informational frictions matter, the overall market climate and investor mood carry greater weight in driving IPO waves.

Interest rate and macroeconomic conditions further modulate the attractiveness of equity issuance. *Monetary policy* and the cost of capital are closely watched by IPO-ready firms. In a low

---

[33] Mandelker, G., & Raviv, A. (1977). Investment banking: An economic analysis of optimal underwriting contracts. *Journal of Finance, 32*(3), 683–694.
[34] Lowry, M. B. (2003). Why does IPO volume fluctuate so much? *Journal of Financial Economics, 67*, 3–40.

interest rate environment, equities become more appealing (both to issuers and investors) as the opportunity cost of investing in stocks falls and abundant liquidity seeks higher returns. Historically, *accommodative monetary conditions have corresponded with upticks in IPO activity.* For example, following the post-2008 recession, the combination of low interest rates and quantitative easing fueled a rebound in global IPO markets: by 2014-2015, IPO volumes surged again, buoyed by easy financing conditions and strong equity rallies. Conversely, rising interest rates or tightening liquidity can cool the IPO market by making investors more risk-averse and increasing firms' cost of equity. Broader macroeconomic indicators like GDP growth, unemployment, and corporate earnings growth also feed into investor sentiment. Generally, economic expansions see more IPOs; firms are growing and need capital, and investors are optimistic, whereas recessions or financial crises coincide with IPO droughts.

1.4.3 Illustrative Periods: Booms, Busts, and Recoveries

Several notable periods illustrate how market conditions swing the fortunes of IPO markets. The late 1990s *dot-com boom* is a quintessential case of a market-driven IPO surge. During 1999-2000, equity markets, particularly technology stocks, soared to unprecedented valuations amid euphoric sentiment. This environment produced an *IPO frenzy*, with record volumes of offerings, especially by internet and tech firms, and extraordinary first-day returns on many IPOs. The European IPO markets partook in this frenzy; exchanges launched new high-growth segments (such as Germany's Neuer Markt) to host the influx of tech listings. However, when the dot-com bubble collapsed in 2000–2001, market sentiment reversed sharply. Stock indices plummeted and investor appetite for new issues evaporated almost overnight. The result was a severe contraction in IPO activity: many planned offerings were shelved, and the number of IPOs in Europe and the US dropped dramatically in the early 2000s. This whipsaw from boom to bust underlines that *bullish market sentiment can rapidly turn, stranding issuers who miss the window.* As one study notes, the early-2000s IPO market was highly unusual and its collapse left a lasting impact on issuance patterns. In the aftermath, regulators and investors grew more cautious, reinforcing how a sentiment-driven boom can sow the seeds of a prolonged cold market.

The *2008 Global Financial Crisis* provides another stark example. Leading up to 2007, many markets were reasonably stable and IPO activity was steady, but the eruption of the financial crisis brought volatility and fear to extreme levels. Equity indices in Europe and the U.S. fell by over 40% at the crisis peak, and the IPO window effectively shut, very few companies dared to go public amid such turmoil. Empirical data confirm that IPO volume suffered a dramatic decline. For instance, in Europe's major markets, IPO issuance dropped to a tiny fraction of previous levels during 2008-2010, reflecting the combination of collapsed valuations and heightened uncertainty. One analysis attributes the low volume of IPOs from 2008 to 2011 at least partly to the lower equity valuations following the 2008 panic and the subsequent Eurozone debt crisis (Ritter et al., 2013)[35].

---

[35] Ritter, J. R., Signori, A., & Vismara, S. (2013). Economies of scope and IPO activity in Europe. In M. Levis & S. Vismara (Eds.), *Handbook of research on IPOs* (pp. 11–34). Edward Elgar Publishing.

In other words, the poor market conditions made equity financing so unattractive (or unattainable) that many firms either waited for recovery or sought alternative funding (such as M&A or private equity). It was only when market conditions improved – stock prices recovering and volatility abating – that IPO activity revived in the early 2010s. This pattern underscores a general principle: *major market downturns and spikes in risk aversion lead to IPO droughts*, whereas sustained rallies and stable conditions revive issuance.

More recently, the *COVID-19 pandemic* shock in early 2020 tested IPO markets worldwide. Initially, the pandemic caused a sudden spike in volatility (e.g. the VIX index reached record highs in March 2020) and a steep drop in stock prices. In Europe, the IPO market became "virtually inactive" for about two months at the height of the uncertainty. Companies that had planned spring 2020 listings largely postponed them, waiting for calmer conditions. By mid-2020, however, unprecedented monetary and fiscal support had stabilized markets, and a strong recovery rally was underway. The post-COVID rebound in market sentiment, coupled with new trends like remote work technology and retail investor participation, led to a flurry of IPOs by late 2020 and a record-breaking IPO boom in 2021. Globally recognized firms (such as Airbnb, DoorDash, etc.) launched highly anticipated IPOs during this period, and European exchanges likewise saw a surge in offerings, particularly in tech, healthcare, and renewable energy sectors. Analysts noted that this 2020-2021 IPO wave was driven by abundant liquidity, low interest rates, and investors' voracious appetite for growth stocks. The episode serves as a vivid illustration of how quickly market conditions can swing: from a sudden freeze in IPO activity during the crisis to an exuberant boom as confidence returned. It also highlights the role of *exogenous shocks* and policy responses in shaping the IPO environment.

1.4.4 Impact on IPO Pricing and Success

Market conditions not only influence the *volume and timing* of IPOs, but also their *pricing and likelihood of success*. In bullish markets with strong demand, issuers can set higher offering prices and still see their shares enthusiastically absorbed by investors (often yielding high initial returns). By contrast, in weak or volatile markets, firms face pressure to *offer deeper discounts* (greater underpricing) to attract cautious investors, or else accept a lower valuation. However, as noted earlier, extremely volatile conditions may simply dissuade firms from even attempting an IPO rather than result in proportionately lower pricing. Underwriters and issuers weigh the reputational and financial risks of a failed offering, and often opt to *delay or withdraw the IPO* if market sentiment deteriorates during the book-building process. In fact, the option to withdraw is an integral feature of the IPO process that connects market conditions to outcomes. Empirical studies find that *poor market returns and high uncertainty correlate with higher IPO withdrawal rates*; for instance, a spike in volatility in late 1998 (during the Russian crisis) saw dozens of U.S. IPO candidates pull their offerings at the last minute. Moreover, evidence indicates that *"management's decision to abort an IPO appears most affected by price uncertainty rather than price levels"*. In other words, a sharp increase in market volatility (even if stock index levels have not fallen drastically) can scare off issuers due to the unpredictable pricing outcomes. This finding underscores how sensitive the success of an equity issuance is to market stability, even a fundamentally strong firm may fail to attract investors if the market backdrop is jittery.

1.5 IPO Withdrawal as a Strategic Response to Volatility

In the previous section, we discussed how prevailing market conditions can influence a firm's decision to issue equity. An important corollary is that when market conditions turn unfavourable, particularly when volatility spikes, firms often respond by *withdrawing* or postponing their planned IPOs. An IPO withdrawal refers to the decision to cancel or delay a public offering after initial filing but before the shares are listed. This section examines IPO withdrawal as a strategic response to volatile markets, drawing on theoretical frameworks and empirical evidence. High volatility environments tend to exacerbate pricing uncertainty and investor risk aversion, making it less attractive for firms to proceed with equity issuance. In practice, the majority of companies that withdraw IPOs cite *unfavourable market conditions* as the primary reason. This behaviour underscores the influence of volatility on financing decisions and reflects a form of market timing strategy by issuers.

1.5.1 Theoretical Perspectives

Several theoretical frameworks shed light on why firms may strategically withdraw IPOs in volatile markets:

- *Real Options Theory:* In corporate finance, managers often treat investment timing decisions as real options. Initiating an IPO can be viewed as analogous to exercising an option to "invest" in public ownership, and volatile conditions increase the value of waiting. When uncertainty (volatility) is high, delaying the offering can be optimal because the firm retains the option to go public later under more favorable conditions. In other words, a volatile market increases the *option value* of postponing the IPO. *Benveniste and Busaba (1997)*[36] and others formalize the idea that the ability to withdraw during the book-building process is essentially a real option for the issuer. This option is more likely to be exercised when market volatility makes the outcome of the offering highly uncertain. Empirical evidence supports this view: one study finds that greater uncertainty significantly increases the likelihood of IPO withdrawal and the value of waiting for volatility to subside.
- *Market Timing Theory:* Market timing arguments suggest that firms seek to issue equity when market conditions are favorable (e.g. high valuations, low volatility) and avoid issuing

[36] Busaba, W. Y., & Benveniste, L. M. (1997). Bookbuilding versus fixed price: An analysis of competing strategies for marketing IPOs. *Journal of Financial and Quantitative Analysis*, *32*(4), 383–403.

when conditions are adverse (low valuations, high volatility). Volatile markets are often "cold" issuance periods in which investor sentiment is poor and IPOs are likely to be undervalued. Firms anticipating this will strategically delay offerings. Chemmanur and Fulghieri (1999) and others have noted that IPO volume tends to concentrate in "hot" markets with bullish conditions and lower uncertainty, whereas during tumultuous periods IPO activity dries up. Withdrawing an IPO during a volatility spike can thus be seen as rational market timing – the firm waits for a window of opportunity when pricing conditions improve. This behaviour aligns with broader evidence that firms practice market timing in financing decisions (e.g. issuing equity when market valuations are high and holding off when they are low).

- *Information Asymmetry and Valuation Uncertainty:* High volatility often reflects greater uncertainty about fundamental values, which can exacerbate information asymmetries between insiders and investors. In an IPO, insiders typically have better information about the firm's true value than outside investors. According to classic adverse selection theory, companies will avoid issuing equity if they believe the market cannot accurately value them and might underprice the shares. In turbulent markets, even high-quality firms fear being undervalued due to noisy signals and cautious investor sentiment. Withdrawing the IPO can thus be a way to avoid selling equity at a price that insiders consider too low. This is consistent with the notion that good firms prefer to wait rather than accept a "lemon's price" during periods of high uncertainty. Empirical research supports this logic: for instance, *Helbing et al.*[37] find that successful IPOs tend to occur when expected market volatility (VIX) is low, whereas volatile periods see more offerings pulled, suggesting that issuers hold off when asymmetric information and valuation risk are elevated. The IPO book-building process itself is a mechanism to reduce information asymmetry by gauging investor interest; if that process reveals unexpectedly low demand (a potential signal that investors have valuation concerns), the issuer may interpret it as an indication of mispricing and choose to withdraw.
- *Signalling Theory:* The decision to withdraw or proceed with an IPO can also send signals to the market. Launching an IPO in adverse conditions could signal desperation for capital, whereas a willingness to withdraw may signal that the firm has other financing options or confidence in its longer-term value. Some signalling models of IPOs (e.g. *Allen and Faulhaber, 1989; Welch, 1989*) suggest that high-quality firms will not sell shares at a depressed price; they might withdraw and wait, implicitly signalling that they believe the firm is worth more than what the market is currently willing to pay. However, signalling effects cut both ways. An IPO withdrawal might alarm some investors by raising questions about why the offering failed. If a withdrawal is interpreted as a sign of weak investor demand or internal problems, it could stigmatize the firm in the short run. Empirical evidence is mixed on this point. On one hand, research by *Boeh and Dunbar (2013)*[38] notes that an IPO withdrawal is not necessarily negative for the firm's prospects if an alternative opportunity (like a private sale) exists, in such cases, withdrawing can be viewed as

[37] Helbing, P., Lucey, B. M., & Vigne, S. (2018). The determinants of IPO withdrawal: Evidence from Europe. *Journal of Corporate Finance.*
[38] Boeh, K. K., & Dunbar, C. G. (2013). *Post IPO withdrawal outcomes.*

positive, strategic choice rather than a failure. On the other hand, studies have found that firms which return for a second IPO attempt often do so at reduced valuations, consistent with the idea that the prior withdrawal sent a cautionary signal to the market (e.g. *Dunbar & Foerster, 2008*[39]). Managers must therefore consider the signal their withdrawal sends: if clearly attributed to external volatility, it may be seen as prudent timing, but if not, investors may infer firm-specific weaknesses.

1.5.2 Strategic Rationale for Withdrawing IPOs in Volatile Markets

Beyond theory, there are practical reasons why IPO withdrawal can be a sound strategic response to market volatility:

- *Pricing Uncertainty and Value Preservation:* Volatile markets make it extremely difficult to price an IPO accurately. Rapid swings in market conditions can widen the gap between the price investors are willing to pay and the price the firm's owners expect based on fundamentals. This creates a risk of severe underpricing (selling shares too cheaply and "leaving money on the table") or, conversely, failing to attract enough demand at the desired price. Rather than proceeding with an offering that could misvalue the firm, companies may choose to withdraw and preserve value for a later date. Studies show that withdrawn IPOs often occur when recent market returns are poor or volatility is high, situations in which the *fair value* of the company is particularly hard to gauge. By waiting for volatility to subside, issuers hope to debut at a price that better reflects their true worth. This rationale is in line with market timing and is frequently cited in practice: for example, *Mikkelson and Partch (1988)*[40] document that the vast majority of withdrawn equity offerings in their sample were attributed to "unfavourable market conditions" or a low stock price environment. Withdrawing thus protects the company from a potentially adverse pricing outcome in a turbulent market.
- *Reputational Risk Management:* Companies and their underwriters are acutely aware of the reputational stakes of an IPO. A failed or poorly received IPO (for instance, one that has to cut its offer price dramatically or that trades down significantly after listing) can tarnish the firm's reputation in capital markets and damage the underwriter's track record. In volatile conditions, the probability of such an outcome is higher, since investor sentiment can swing unpredictably. By withdrawing the IPO, the issuer can avoid the spectacle of a failed offering. This is often seen as the lesser of two evils: a quiet withdrawal draws less public attention than a spectacular IPO failure. Empirically, it is observed that firms that withdraw often face negative media coverage around the time of the withdrawal, but this short-term scrutiny may be preferable to the sustained reputational hit of a disastrous IPO. From the underwriter's perspective, withdrawal can be a prudent choice as well. Top investment banks will counsel clients to postpone if the order book is weak, rather than

---

[39] Dunbar, C. G., & Foerster, S. R. (2008). Second time lucky? Withdrawn IPOs that return to the market. *Journal of Financial Economics.*
[40] Mikkelson, W. H., & Partch, M. M. (1988). Withdrawn security offerings. *Journal of Financial and Quantitative Analysis, 23*(2), 119–133.

push a deal that could flop. Notably, research finds that companies who withdraw are not necessarily small or low-quality; in fact, they tend to hire underwriters just as reputable as those used by successful IPOs. This suggests that even high-reputation intermediaries and mature issuers opt to withdraw in the face of volatile, adverse market feedback. The ability to abort the issue is thus part of maintaining long-term reputational capital for both issuers and underwriters.

- *Book-Building and Investor Feedback*: The modern IPO process (especially in the US and Europe) gives issuers a mechanism to test the waters before committing to an offering. Through book-building, underwriters gather indications of interest from institutional investors, which provides a sense of the demand curve for the IPO. If this feedback indicates that investors would only be willing to buy at a price significantly below the initial range, or if the overall market mood has soured during the book-building period, the issuer can exercise the option to withdraw. In fact, the option to withdraw is embedded in most underwriting contracts: issuers typically have the right to pull an offering, usually with minimal financial penalty (often just reimbursing the underwriter's out-of-pocket expenses). This flexibility is crucial in volatile times. It means that firms are not forced to complete an IPO come what may; they have a built-in escape hatch if conditions deteriorate. Underwriter discretion plays a role here as well. Underwriters may expand the syndicate or adjust marketing efforts to gauge broader investor sentiment (as noted by Corwin & Schultz, 2005[41]), and if those efforts still yield weak demand, they will advise cancellation. The book-building process, therefore, acts as an information discovery and risk management tool: it allows the issuer to conditionalize the IPO on receiving a favorable enough reception. If not, *no deal* is a viable outcome. This mechanism is one reason why IPO withdrawal rates tend to rise during volatile periods – the process itself flags the heightened risk and provides a structured way to pull back.

- *Regulatory and Timing Considerations:* Various regulatory regimes provide windows or time limits within which an IPO must be completed after regulatory approval. If market volatility is high, companies sometimes choose to formally withdraw their registration and re-file later, rather than rushing to market under poor conditions. In the U.S., for example, a firm that has filed with the SEC can submit a withdrawal request (Form RW) and face no punitive action, preserving the possibility to re-launch the IPO when the window reopens. European and other markets likewise allow cancellation of an offering before listing, in many cases with the ability to reuse the prospectus after updates. These regulatory provisions effectively acknowledge that market timing is a legitimate concern. They enable issuers to *wait out* short-term volatility spikes. From a strategic standpoint, management will weigh the costs of delaying (such as additional administrative expenses or the risk of missing a transient favorable window) against the risks of proceeding. The opportunity to withdraw

[41] Corwin, S. A., & Schultz, P. (2005). The role of IPO underwriting syndicates: Pricing, information production, and underwriter competition. *Journal of Finance*, 60(1), 443–486.

thus provides critical flexibility in volatile environments, and wise issuers use it when the downside of pushing forward outweighs the benefits of waiting.

1.5.3 Empirical Patterns of IPO Withdrawal

The strategic behaviour described above is reflected in empirical patterns observed across IPO markets. IPO withdrawal is a common feature of equity markets worldwide, though its prevalence varies over time and across regions. In the United States, a substantial fraction of IPO filings do not result in completed offerings, roughly 30% of IPO filings during the early 2000s were eventually withdrawn; in Europe, withdrawal rates have been somewhat lower, around 12% on average in a 2001-2015 sample, but still significant (see Helbing et al., 2018).

These averages conceal considerable time variation driven by market volatility. Withdrawal rates tend to spike in turbulent periods and recede in stable or bullish periods. During the Global Financial Crisis and its aftermath, IPO markets virtually shut down: there were very few new offerings, and many of those that were attempted had to be pulled. In 2009, amid extreme volatility, there were almost no IPOs in many major markets (only 18 IPO filings in all of Western Europe that year) and consequently very few completions. A couple of years later, the Eurozone debt crisis of 2011 again rattled markets; Europe's IPO withdrawal rate surged to about 22% that year, the highest in the sample. Conversely, in calmer periods such as 2003-2006, withdrawal rates were much lower (in some years under 5%) as strong market conditions facilitated IPO success. This inverse relationship between market volatility and IPO completion is well documented. Schill (2004), for instance, finds that increases in market volatility are strongly associated with increases in the fraction of IPOs withdrawn, even after controlling for other factors. In addition to macro volatility, market returns and investor sentiment also play a role: IPO withdrawals are more likely when recent stock market returns are negative or when there is broad uncertainty, reinforcing the idea that firms try to avoid issuing into a falling or volatile market.

Cross-sectional studies have identified certain characteristics of withdrawn IPOs. Interestingly, firms that withdraw are not necessarily fundamentally weaker than those that complete. *Busaba et al. (2001)* report that withdrawn IPO candidates in the U.S. were, on average, similar or larger in size and profitability compared to those that went public, and they often had venture capital backing and reputable underwriters, indicating that many were high-quality firms. However, these firms tended to have higher leverage and were more likely to have intended to use IPO proceeds to repay debt, suggesting a greater sensitivity to pricing (a heavily indebted firm might withdraw if the pricing isn't sufficient to meet its debt reduction goals). Consistent with information asymmetry ideas, withdrawn IPOs also coincide with situations of high ex-ante uncertainty, for example, companies in high-tech or innovative sectors may pull offerings if investors appear unable to agree on valuation. In Europe, research by *Helbing, Lucey, and Vigne* finds that private equity or venture-backed firms are overrepresented among withdrawals, often because these firms have viable alternatives such as trade sales (mergers & acquisitions). These firms can pursue a *"dual-track" strategy*: they file for an IPO and simultaneously engage in negotiations with potential acquirers. If the IPO valuation comes in too low or market volatility makes the public route unattractive, they pivot to a sale. Indeed, many withdrawn IPOs end up as successful acquisitions, sometimes at valuations higher than what the IPO would likely have achieved. Such outcomes illustrate that an IPO withdrawal can be a strategic rerouting towards a different objective (getting bought out or waiting for a better IPO window) rather than an outright failure.

1.5.4 Consequences and Trade-Offs of Withdrawing an IPO

Choosing to withdraw an IPO carries several consequences, both positive and negative, that management must carefully weigh:

- *Opportunity Cost and Financing Delay:* The most immediate consequence is the delay in capital raising. If the IPO is withdrawn, the firm does not receive the expected equity capital on the original timeline. For companies that needed the funds for expansion, R&D, debt reduction, or other projects, this delay can be costly. Some firms may have to turn to alternative financing sources (e.g. private placements, venture capital top-ups, bank loans) to bridge the gap. In certain cases, a withdrawn IPO is only a temporary setback and the firm refiles when conditions improve; in others, the firm might shelve growth plans or seek a strategic partner instead. The opportunity cost of waiting can be high if market conditions take long to improve or if competitors move ahead in the interim. However, this must be balanced against the potential cost of having gone public under unfavourable conditions (e.g. raising much less money than anticipated or diluting existing owners at a low price).
- *Probability of Re-filing and Future Success:* Empirical studies show that many companies that withdraw an IPO do not return to the public markets for a long time, if ever. For example, *Dunbar & Foerster (2008)* and subsequent studies found that issuers who withdraw are unlikely to attempt a second IPO in the near term, and if they do, the likelihood of success can be lower. Those that do re-file often face a credibility hurdle. They must demonstrate that the factors which led to the previous withdrawal (be it market conditions or firm issues) have resolved. There is evidence that when withdrawn firms eventually do go public, they tend to do so at lower offer valuations than originally hoped. This implies that withdrawing can have a lasting impact on the terms the firm can achieve, possibly because investors demand a discount to compensate for the uncertainty created by the prior withdrawal. On the flip side, firms with strong fundamentals that withdrew due to purely external reasons (e.g. a market crash) might see no permanent damage; they might re-list successfully when the window reopens. Each case is unique, but statistically the delay between IPO attempts can be substantial, and some firms abandon the IPO path altogether in favour of remaining private or selling the company.
- *Signalling and Investor Perceptions:* As discussed, an IPO withdrawal can send a negative signal to the market regarding investor interest or the firm's prospects. If not managed properly, it can create a stigma: investors might wonder if the withdrawal was due to undisclosed problems. To mitigate this, companies typically attribute withdrawals to market volatility or other external factors (often truthfully so). During broad market downturns, a withdrawal likely will not be held against the firm, investors understand that many offerings are pulled in such environments. However, in more stable times, a solitary withdrawal is more suspect. There can also be an impact on the firm's relationships with underwriters and investors. Underwriters, while understanding of market-driven withdrawals, will keep

records, a company that has pulled an offering may face tougher scrutiny or more conservative pricing in a future attempt. On the investor side, those who indicated interest in the IPO may be disappointed or cautious the next time around. In sum, withdrawing is *not* a costless choice: it can alter the firm's reputation in capital markets. Nonetheless, when the alternative is a catastrophic IPO, withdrawing is often the wiser long-term decision, and market participants recognize that. For example, *Boeh and Dunbar (2013)* emphasize that if a firm has a "superior option" (like a lucrative buyout offer or the ability to wait for a better market), then withdrawing from the IPO can ultimately be seen as a positive outcome for shareholders. The key for management is to communicate clearly and maintain credibility, so that a withdrawal is viewed in context as strategic patience rather than a sign of internal failure.

- *Dual-Track Outcomes and Shareholder Value:* Especially for venture-backed companies, an IPO withdrawal may simply mark a pivot to another liquidity event. Studies focusing on post-IPO withdrawal outcomes in Europe document a high incidence of M&A transactions shortly after withdrawal, suggesting that many firms are taken over at valuations that can reward early investors and founders despite the IPO not happening. In such cases, withdrawing the IPO can actually maximize shareholder value, the firm avoids a weak IPO and secures a better price through acquisition. Alternatively, the firm may remain private and continue to grow, waiting for an eventual IPO under more auspicious circumstances. There is also evidence that the "death" rate (bankruptcy or dissolution) for withdrawn IPO firms is not unusually high in the short run, implying that most withdrawn firms find a way to survive either through private funding or being absorbed into another company. Thus, the outcome of an IPO withdrawal is often a redirection rather than an end. Shareholders and boards will consider these trade-offs when approving the decision to pull an offering: is there a reasonable expectation of achieving a better outcome by waiting or pursuing a sale? If yes, withdrawal can be justified as part of a savvy corporate strategy.

IPO withdrawal serves as a critical *safety valve* in corporate financing decisions under volatility. It allows firms to navigate uncertain markets without committing to potentially unfavourable transactions. This strategic flexibility is supported by real options reasoning (delay when risk is high), aligns with market timing behaviour, and is facilitated by the IPO process design (book-building and regulatory provisions). However, it comes with its own costs and considerations, including delayed financing and potential signalling effects. The decision to withdraw is therefore a nuanced one, balancing short-term market conditions against long-term corporate finance objectives.

Given the importance of volatility in triggering IPO withdrawals and other financing decisions, a natural question is how companies and underwriters can anticipate or manage this volatility. Accurately forecasting market volatility becomes crucial in planning the timing of equity issues.

1.6 Forecasting Market Conditions: The Role of GARCH and Other Volatility Models

Building on the discussion of IPO withdrawal as a response to volatility, we now examine how firms can forecast market volatility as part of their financing decisions. Corporate finance decisions, especially the timing of an IPO, are inherently forward-looking, meaning that expected market conditions matter as much as current conditions. High volatility often deters equity issuance, but if firms can anticipate volatile periods, they may adjust their timing or pricing strategy accordingly. Prior studies show that IPO activity tends to be higher during "calm" or favorable markets (low volatility), whereas uncertain markets see more hesitation. At the same time, the option to withdraw or delay an IPO gives issuers flexibility, a flexibility that is *more valuable when volatility is high*. This real-options perspective underscores why volatility forecasts, not just observed volatility levels, are crucial: by predicting volatility, firms can better decide whether to proceed with an offering or wait for a more opportune window.

1.6.1 Approaches to Volatility Forecasting in Finance

Volatility forecasting has become a central topic in finance, and numerous methods have been developed to predict future market volatility. Broadly, the approaches can be grouped into three categories: (1) historical volatility models, (2) implied volatility from options, and (3) econometric models such as GARCH and its variants. Each approach provides a different lens on future volatility, with its own strengths and weaknesses, particularly relevant in the IPO context (where firms are concerned with near-term market uncertainty around the issuance date).

1. *Historical Volatility Models (Moving Averages and EWMA):* The simplest approach is to use past price data to estimate volatility and assume it will persist. For example, a *rolling window* calculation of volatility (e.g. using the standard deviation of daily returns over the past 60 days) provides a quick estimate of current volatility. An alternative is the Exponentially Weighted Moving Average (EWMA) method popularized by J.P. Morgan's RiskMetrics, which gives more weight to recent observations.

   The RiskMetrics formula updates variance as

   $$h_t^2 = \lambda h_{t-1}^2 + (1-\lambda) r_{t-1}^2$$

   whereas the decay factor λ is typically set to 0.94. These historical models are easy to implement and react to changes in volatility to some degree. However, their forecasts are purely backward-looking and slow to adapt if the volatility regime shifts suddenly. A fixed rolling window may either lag true volatility (if the window is long) or produce noisy forecasts (if the window is too short). Likewise, the EWMA assumes a constant decay rate and cannot account for structural breaks or asymmetric reactions to news. Importantly, historical approaches do not incorporate any market expectations beyond what is in the past data, which can be a limitation in fast-changing conditions (as often seen leading up to IPOs).

2. *Implied Volatility from Options:* A more forward-looking approach comes from the prices of derivative contracts. Implied volatility is the volatility that equates an option's market price with a theoretical pricing model (typically Black–Scholes). It effectively represents investors' collective expectations of future volatility. For example, the CBOE's VIX index (derived from S&P 500 options) is often called the market's "fear gauge" and is interpreted as the expected stock market volatility over the next 30 days. Implied volatility is a market-based forecast, aggregating information from many market participants. In many studies it has been found to contain information beyond what historical returns alone provide – in fact, implied volatilities often outperform time-series models in predicting future volatility. This makes intuitive sense: option prices respond quickly to new information and risk sentiment, so the implied volatilities embed real-time forward-looking adjustments. However, there are notable caveats to using implied volatility. First, implied volatility can be a biased predictor of realized volatility, for instance, it often *overestimates* future volatility on average, due to volatility risk premia and other factors (implied vols tend to be higher than realized vols, especially in calm periods). Poon and Granger (2005)[42] document that implied vol forecasts are usually biased but still useful, as they may consistently over-predict low-volatility outcomes. Second, for many firms or specific contexts, suitable option data may not be available. In the context of IPOs, the firm going public typically has no traded options before the IPO, this means that implied volatility is not directly observable for that stock. Instead, practitioners might look at implied volatilities on a market index or a sector ETF as a proxy for market conditions. Indeed, recent IPO research uses indices like the VIX as a proxy for market volatility expectations when analysing issuance decisions. The strength of implied volatility forecasts lies in being forward-looking and quickly responsive to new information, but their accuracy can be influenced by market microstructure noise, liquidity, and the assumption of the pricing model.

3. *Econometric Volatility Models (ARCH/GARCH Family):* A large class of models directly estimate and forecast volatility using past time-series data, while allowing volatility to be time-varying (stochastic). The work of Engle (1982)[43] on Autoregressive Conditional Heteroskedasticity (ARCH) and Bollerslev (1986)[44] on Generalized ARCH (GARCH) introduced models that capture the stylized fact of *volatility clustering*, the tendency for volatility to come in waves (periods of high volatility followed by high, and low followed by

[42] Granger, C., & Poon, S.-H. (2005). Practical issues in forecasting volatility. *Financial Analysts Journal, 61*(1).
[43] Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, *50*(4), 987–1007.
[44] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327.

low). A basic GARCH(1,1) model forecasts tomorrow's variance as a function of yesterday's squared return (news/shock) and yesterday's variance, effectively saying that *volatility tomorrow depends on recent volatility and recent surprises.* This framework has proven remarkably successful in forecasting contexts; indeed, the GARCH(1,1) often serves as a robust benchmark. Hansen and Lunde (2005)[45] famously found no evidence that more complex models could consistently outperform a simple GARCH(1,1) in out-of-sample volatility prediction. The strengths of GARCH models include their ability to *adapt to changing volatility* (unlike a fixed historical window) and to produce multi-step forecasts in a coherent way under the model. They are relatively parsimonious yet capture the persistence in volatility typically observed in markets. That said, basic GARCH assumes volatility responds symmetrically to past shocks, an assumption at odds with the observed "leverage effect," where negative returns often cause larger volatility jumps than positive returns of equal magnitude. To address this, researchers have developed GARCH variants that incorporate asymmetry. Two prominent examples are the *Exponential GARCH (EGARCH)* by Nelson (1991)[46] and the *Glosten–Jagannathan–Runkle GARCH (GJR-GARCH)* by Glosten *et al.* (1993)[47]. EGARCH models the logarithm of variance and includes a term to allow negative surprises (bad news) to have a different impact than positive surprises. This way, EGARCH can capture the fact that market downturns tend to cause disproportionately high volatility spikes. Similarly, GJR-GARCH (also known as a Threshold GARCH) adds a dummy term that "kicks in" when returns are negative, effectively permitting volatility to respond more to downside shocks. These extensions significantly improve forecasting in markets where the leverage effect is strong (e.g. equity markets where bad news leads to higher uncertainty). Overall, the GARCH family provides a flexible toolkit for forecasting: it is possible to choose specifications to handle long-memory (e.g. FIGARCH for slowly decaying volatility), structural breaks, or other features as needed. The weakness, however, is that these models require careful specification and estimation. They rely solely on historical prices (thus *not* incorporating outside forward-looking information like option markets do) and can be misled by regime changes or novel shocks (unless supplemented with additional inputs). Still, in practice, a well-specified GARCH or its variant often yields a solid volatility forecast and serves as a base for more complex models or forecast combinations.

1.6.2 Strengths and Weaknesses in the IPO Timing Context

In choosing a forecasting approach for IPO decision-making, one must consider the trade-offs. *Historical models* (like rolling stdev or EWMA) are simple and transparent, which might appeal to practitioners and can be quickly computed even when data is limited. However, their

---

[45] Lunde, A., & Hansen, P. R. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics, 20*(7), 873–889.
[46] Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, *59*(2), 347–370.
[47] Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance,* 48, 1779-1801.

reactive nature means they might not warn an IPO issuer of an upcoming volatility surge, they would only register it once it has begun. *Implied volatility* is appealing because it directly reflects market expectations. If, for example, the implied volatility of a broad market index is rising, an IPO issuer might take that as a warning signal of turbulent times ahead. The downside is that implied metrics can sometimes be influenced by factors other than pure volatility expectation, such as demand/supply imbalances in the options market or risk premia; and as noted, they are not firm-specific. *GARCH-type models* offer a middle ground: they use statistical patterns in historical data to project forward, often giving reliable short-term forecasts and capturing the volatility persistence that simpler models miss. A GARCH forecast could, for instance, quantify the risk that volatility remains elevated in the weeks after a market shock, information very relevant to an IPO's after-market stability. Yet, if an unprecedented event is on the horizon (something not reflected in past data), a pure GARCH model will not foresee it. In summary, no single method dominates in all circumstances: each has advantages that might be complementary. In fact, research suggests that combining forecasts (e.g. using both GARCH and implied vol) can improve accuracy. Kambouroudis *et al.* (2016)[48] find that while implied volatility alone may perform worse than some model-based forecasts, a hybrid model blending an asymmetric GARCH with implied and realized volatility information performs best. This insight will guide the empirical strategy of this thesis, as we next discuss.

1.6.3 Realized Volatility and Forecast Evaluation Methods

To evaluate which volatility forecasting approach is most effective, we need a reliable benchmark of true volatility to compare forecasts against. Volatility itself is latent (not directly observable), but we can construct ex-post measures of realized volatility from high-frequency or daily data. *Realized volatility* refers to the actual volatility realized over a period, typically measured by the variance or standard deviation of asset returns within that period. Advances in financial econometrics have shown that using high-frequency (intra-daily) data can yield very accurate realized volatility measures. For example, summing the squared 5-minute returns of a stock over a day gives a close estimate of that day's actual variance. Such high-frequency-based realized volatilities provide valuable insight into the true latent volatility, in practice, for evaluation, one might use daily realized volatility computed from intraday data (if available) or, alternatively, the sum of squared daily returns over a month as a proxy for that month's volatility. The accuracy of forecasts can then be assessed by comparing the forecasted volatility for a period to the realized volatility that actually occurred in that period.

When comparing volatility forecasts, it is important to use appropriate *loss functions* that quantify forecast errors. A common choice is the *Mean Squared Error (MSE)*, which simply squares the difference between the forecasted variance and the realized variance. MSE is intuitive but can be sensitive to extreme errors, and in volatility forecasting it may not always be the most statistically powerful metric. An alternative that is widely recommended in volatility forecast evaluation is the

---

[48] Kambouroudis, D. S., McMillan, D. G., & Tsakou, K. (2016). Forecasting stock return volatility: A comparison of GARCH, implied volatility, and realized volatility models. *Journal of Futures Markets*, 36(12), 1127–1163.

so-called *QLIKE* (Quadratic Likelihood) loss function. The QLIKE loss is a type of scaled log-error that has desirable robustness properties – notably, it remains valid even if the realized volatility is a noisy proxy for true volatility. Patton and Sheppard (2009)[49] argue that using QLIKE can lead to more reliable comparisons between models, and they find that forecast evaluation tests (like Diebold-Mariano) often have higher power under QLIKE than under MSE. In this thesis, we will employ both MSE and QLIKE to ensure that results are not sensitive to the choice of loss function, while recognizing the benefits of QLIKE in volatility contexts.

To formally test whether one forecasting model is significantly better than another, we utilize forecast evaluation tests. The most prominent is the *Diebold–Mariano (DM) test* for equal predictive accuracy (Diebold & Mariano, 1995)[50]. The DM test considers the loss differentials of two forecasts (e.g. the difference in squared errors) and assesses whether the average difference is zero or significantly positive/negative. Essentially, it tells us if one model's forecasts are statistically more accurate than another's. For multiple model comparisons, a more recent approach is the *Model Confidence Set (MCS)* procedure of Hansen, Lunde, and Nason (2011)[51]. The MCS is an iterative testing framework that aims to identify a set of models that includes the best model with a given confidence level, by eliminating significantly inferior models. In other words, rather than picking a single "winner," MCS yields a subset of forecasting models that are indistinguishably best at a certain confidence level. This is particularly useful in volatility forecasting studies where many models are compared, we avoid false certainty in ranking and instead report a confidence set of top performers. Employing DM tests and MCS will allow us to rigorously compare forecasts from historical, implied, and GARCH models in our empirical analysis.

1.6.4 Why Volatility Forecasts Matter for Corporate Decisions

It is worth emphasizing why all this matters for corporate financing decisions. Volatility forecasts provide information that current volatility alone cannot, namely, they inform decision-makers about the *anticipated risk environment* during the window in which a financing transaction (like an IPO) will take place. If a company is considering going public, knowing that the market is currently calm is helpful, but knowing whether that calm is likely to continue or if turbulence is on the horizon is even more crucial. IPOs involve processes that unfold over weeks or months (from initial filing to pricing to first trading). Thus, what matters to issuers and underwriters is the *expected market volatility* over the offering period. A spike in volatility can increase the risk of underpricing or even of a failed offering, as investor sentiment may sour or valuations become more uncertain. By using volatility forecasts, firms can attempt to time their IPOs to avoid stormy market conditions, or conversely, seize windows of opportunity when volatility is predicted to stay low. As an example, suppose forecasts indicate a high probability of volatile swings in the next quarter due to political or

---

[49] Patton, A. J., & Sheppard, K. (2009). Evaluating volatility and correlation forecasts. In T. G. Andersen, R. A. Davis, J.-P. Kreiss, & T. Mikosch (Eds.), *Handbook of financial time series*.
[50] Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*.
[51] Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The model confidence set. *Econometrica*, *79*(2), 453-497.

economic events; a firm might choose to postpone its IPO to avoid that period. On the other hand, if forecasts show a stable outlook, the firm gains confidence to proceed. In short, volatility forecasts act as a form of *advance warning system* that can shape strategic decisions such as *whether to withdraw, delay, or go ahead with an IPO*. This forward-looking risk management is in line with the real options view mentioned earlier: since firms have the option to delay an IPO, having a reliable forecast of volatility helps in exercising that option optimally. It is not just the level of volatility but the *change* or trend in volatility that can drive decisions, a market transitioning from low to high volatility can be especially perilous for issuers, even if current volatility is moderate.

2 Empirical Analysis Part I: Volatility Modelling and Forecasting

In Chapter 1 we introduced the motivation for our analysis of European equity markets and volatility, reviewing prior research on financial time series and volatility modelling. Next, we lay the empirical foundations and describe the dataset and the methodological framework used to pursue the objective of answering the following question: "How do forward-looking measures of equity market volatility, derived from historical models, implied volatility, and GARCH-type forecasts, affect the likelihood of IPO withdrawal, the revision of offer prices, and first-day returns in the European IPO market?"
The empirical work is organized in two distinct but interconnected parts, corresponding to Chapters 2 and 3.

The first part, covered in Chapter 2, addresses the first part of that question and hence focuses on volatility itself and how to model and forecast it using multiple methodologies and datasets. To be more precise, we compare different ways of quantifying and predicting equity market volatility, both backward-looking (historical models like EWMA or rolling-window estimators) and forward-looking (implied volatility from options, and parametric forecasts from GARCH-family models). The objective of this first empirical analysis is to assess the predictive power and stability of various volatility forecasting tools. This is essential because volatility, as established in Chapter 1, is one of the most influential market conditions affecting equity issuance decisions. Before we can analyse those financing outcomes, we need a robust, dynamic measure of market uncertainty.

The second part, developed in Chapter 3, builds directly on this foundation and links volatility forecasts to real IPO outcomes. Using a dataset of European IPOs from 2010 to 2024, the study tests how market conditions, proxied by volatility estimates, affect the firm's decision to withdraw the IPO, revise the offer price, or experience higher initial returns (underpricing). These are not separate research directions: the goal of the thesis is to empirically connect *forward-looking market risk* (as seen from volatility forecasts) with *real corporate financing behaviour* (as reflected in IPO outcomes). This two-stage structure allows for a clear analytical flow: first, establish how volatility behaves and can be predicted (Chapter 2), then study how those forecasts influence financing decisions in practice (Chapter 3).

2.1 Data for Volatility Forecasting

This section details the data sources and preparation used to construct volatility forecasts for the Eurozone equity market. We focus on two time series: the *EURO STOXX 50 Index (ticker SX5E)*, a benchmark equity index, and the *V2X Index (also known as VSTOXX)*, a volatility index for the EURO STOXX 50. We describe their economic roles, statistical characteristics, and relevance for capturing Eurozone market conditions, supported by literature. We then justify the sample period (1 January 2005-31 December 2024) in terms of market cycles, volatility regimes, and IPO cycles, noting major structural breaks (e.g. the Global Financial Crisis, Eurozone debt crisis, COVID-19, and energy shock). The data structure is daily: SX5E provides closing, high, and low prices; V2X provides daily closing values. We explain how these price series are converted into daily log returns (with references to financial econometrics) and discuss treatment of missing values and non-trading days. Finally, we perform exploratory data analysis, presenting descriptive statistics (mean, variance, skewness, kurtosis) and discussing stylized facts such as heavy tails, volatility clustering, and the leverage effect, as well as the negative correlation between equity returns and volatility.

2.1.1 Euro STOXX 50 Index (SX5E) – Price Series and Historical Volatility

The SX5E index is a free-float market-capitalization-weighted equity index representing the largest 50 blue-chip companies in the Eurozone. It spans 20 industry *supersectors* and covers the largest and most liquid firms in Europe; it is widely regarded as the leading benchmark for Eurozone equity performance.

This index captures roughly 60% of the market capitalization of the total Euro STOXX broad market index. By construction, SX5E constituents are updated semi-annually to ensure liquidity and representativeness. It is heavily used as an underlying for derivatives (options and futures), ETFs, and benchmarks in Europe. Because it covers multiple countries (e.g. Germany, France, Spain, Italy, etc.) and sectors (financials, industrials, consumer goods, etc.), the SX5E provides a broad measure of Eurozone market conditions. Its composition and weighting scheme (free-float market cap) are documented in STOXX methodology guides.

The dataset includes three variables: daily closing level, intraday high and intraday low for SX5E on each trading day. The closing price is the end-of-day index value, which is used to compute daily returns as the basis for historical volatility. High and low prices denote the day's extreme values and can provide additional information about intraday variability. For instance, the range (high minus low) on a given day is itself an informative volatility proxy: Parkinson (1980)[52] famously showed that the daily high–low range, properly scaled, is an unbiased estimator of the true variance and about five times more efficient than using only close-to-close returns. Although our primary measure of realized volatility will be derived from daily closing prices (e.g. squared returns or rolling window estimates), the inclusion of high/low data allows for range-based volatility estimates to be examined and used for robustness checks. All price data are in index points (EUR-based) and reflect the official SX5E index values published by STOXX, which inherently adjust for corporate actions

[52] Parkinson, M. (1980) 'The extreme value method for estimating the variance of the rate of return', *Journal of Business*, 53(1), pp. 61–65.

of constituents. No additional adjustments (such as for dividends or stock splits) are necessary, as the index methodology maintains consistency over time.
The SX5E price series serves as the foundation for *historical (realized) volatility* calculations in our analysis. In volatility forecasting, realized volatility refers to the ex-post actual volatility of the asset, against which forecast performance is evaluated. Since volatility is latent, we will compute realized volatility measures from the SX5E data, for example, using daily log-return variances or aggregating higher-frequency variations (if available). In our daily context, a common approach is to use the variance of daily returns over a certain window or simply the next day's squared return as a one-day realized volatility. The long sample of SX5E daily prices (over 5,000 trading days in 2005-2024) yields a rich set of return observations for estimating and testing time-series volatility models (e.g. GARCH family models to be introduced in later sections). The choice of a major stock index like SX5E ensures that our volatility forecasts pertain to a broad market portfolio, which is standard in the literature and relevant for investors. Moreover, indices are less prone to idiosyncratic shocks and have well-documented volatility characteristics (such as leverage effects and volatility clustering) that volatility models can capture.

2.1.2 V2X Volatility Index (VSTOXX) - Implied Volatility Measure

The V2X index is the *EURO STOXX 50 Volatility Index*, often referred to by its alias VSTOXX. It is analogous to the VIX index for the U.S. market, providing a *forward-looking* gauge of expected stock market volatility. The V2X is derived from prices of EURO STOXX 50 index options using a model-free methodology (effectively the pricing of a 30-day variance swap) and is quoted in volatility percentage points, annualized. In other words, V2X reflects the market's expectation of the *annualized standard deviation* of SX5E returns over the next 30 calendar days. A level of, say, 20 on the V2X corresponds to an implied volatility of 20% per annum for the coming month. As an *implied volatility* index, V2X encapsulates information from option market participants about future uncertainty. This makes it a natural benchmark for volatility forecasting: it represents the "collective forecast" of market volatility by informed traders and is frequently found to be highly informative about subsequent realized volatility. Numerous studies have shown that implied volatilities like VIX/VSTOXX often outperform purely backward-looking estimates in forecasting future volatility, since they aggregate forward-looking information and risk premia. In our context, the V2X series will serve as the key proxy for market-anticipated volatility against which we can compare model-based forecasts. We may use V2X in two ways: (a) as a benchmark to evaluate the accuracy of our econometric forecasts (i.e. asking whether our models can add value relative to what is implied by the options market), and (b) potentially as an input (explanatory variable) in extended volatility models, given the strong information content of implied volatility noted in prior research. The V2X data extracted from the Bloomberg Terminal provides the daily closing value of the volatility index; this aligns in time with the SX5E closing price (since both are determined at the end of the trading day, with V2X typically computed from end-of-day option prices). Because V2X is based on a liquid options market and maintained by STOXX, the data quality is high, and it already excludes non-trading days (matching the SX5E trading calendar closely).

In sum, SX5E and V2X together capture both the level of the Eurozone equity market and market-implied volatility. Using these indexes allows us to model and forecast volatility at the Eurozone market-wide level rather than relying on a single country or sector.

2.1.3 Sample Period, Frequency and Data Preparation

The chosen sample spans 20 years, covering various market cycles, regimes, and IPO waves. This period includes both tranquil and turbulent times. We capture the strong bull market of the mid-2000s, the 2007-2009 Global Financial Crisis (GFC), the European Sovereign Debt Crisis of 2010-2012 (Greek, Irish, Portuguese debt turmoil), the relatively calm period of 2013-2017, the volatility spike in early 2016 (China and oil shocks), the Brexit episode (2016), the COVID-19 shock in 2020, and the recent inflation/energy crisis (2021-2022). By including 2005 onward, we observe the lead-up to the GFC (the tail end of the housing bubble era) and its aftermath. The inclusion of 2023-24 is also crucial, as these years saw renewed volatility due to geo-political tensions (e.g. Ukraine war) and global inflation, ensuring up-to-date coverage.

Each crisis introduced structural breaks in volatility. For instance, the GFC (Lehman bankruptcy, Sept 2008) caused an unprecedented spike in volatility and correlated global market declines. The VSTOXX and SX5E both exhibited dramatic moves in late 2008. Similarly, the Eurozone debt crisis (2010-2012) caused sustained market stress. For example, Greek debt and potential euro breakup fears in 2011 sent volatility higher and equities lower. The COVID-19 pandemic (March 2020) led to one of the fastest spikes in volatility on record, with VSTOXX surging. More recently, the European energy shock and inflation in 2021-22 produced renewed volatility. We ensure our dataset captures these regimes to allow models to learn from genuine stress events.

The sample also includes typical IPO cycles. For example, many large European IPOs occurred in the 2006-2007 window (pre-GFC), and another IPO boom around 2014-2017. Volatility tends to influence IPO timing and pricing: firms may withdraw IPOs or revise prices when markets are volatile. By covering years with different IPO volumes and market sentiments, we enable testing how forecasted volatility correlates with IPO outcomes. Empirically, IPO volumes fell sharply in crisis years (e.g. 2008, 2020) and rose in calmer periods, making volatility forecasts relevant. Citing IPOs specifically is beyond this data chapter, but the timeline ensures volatility measures align with the IPO dataset and underlying economics.

From an econometric viewpoint, including 2005-2024 enhances sample size, improving estimation precision. It also incorporates possible parameter shifts: for instance, volatility dynamics changed post-2008 (higher baseline volatility, different GARCH parameters). While formal break tests (e.g. Bai-Perron) could identify regime changes, we instead explicitly acknowledge known breakpoints (2008, 2010-12, 2020) when interpreting results. If needed, regime dummies or robust methods can handle these. At minimum, describing these events ensures any volatility models account for structural volatility shifts (e.g. conditional variance intercept changes).

The frequency of the data is daily, which is a standard choice for volatility modelling in academic studies when high-frequency data (intraday prices) are not utilized. Daily frequency offers a balance between capturing the persistence in volatility and limiting noise: volatility clustering is clearly visible at the daily level, and most classic models (e.g. GARCH) are naturally specified for daily return series.

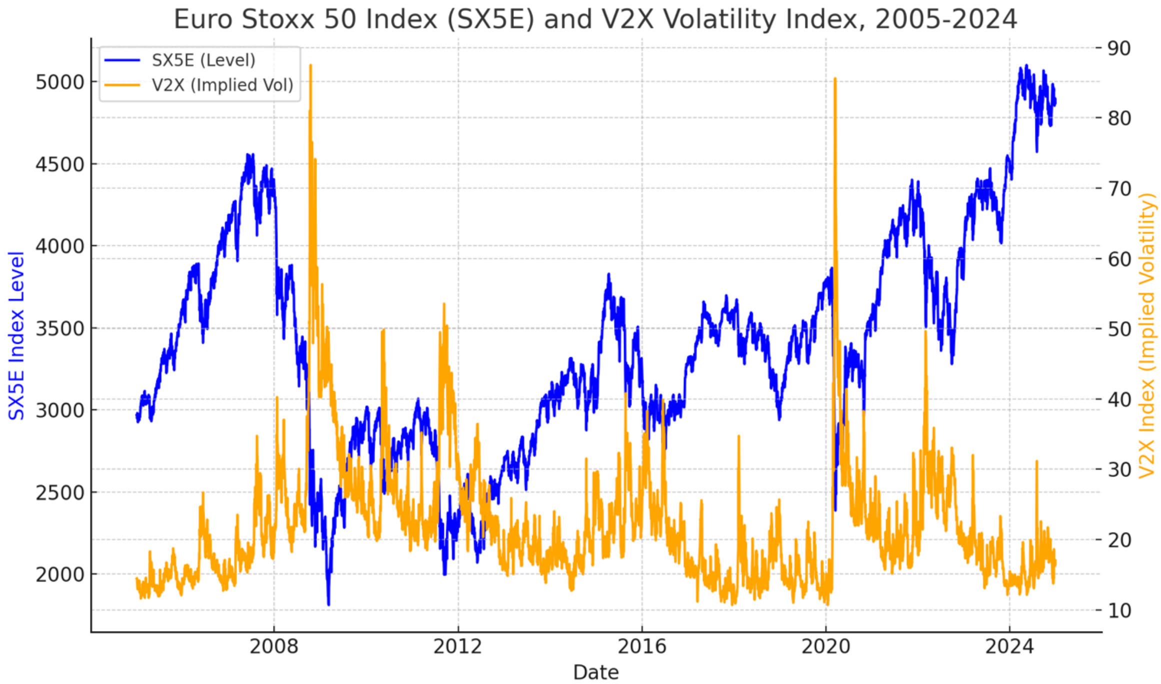

Figure 2.1: Evolution of the Euro STOXX 50 index (SX5E, blue, left axis) and the V2X implied volatility index (orange, right axis) over the sample period 2005–2024. The SX5E index level exhibits several major swings, notably a steep decline during the 2008–2009 financial crisis and another drop during the early 2020 pandemic, followed by recoveries. Correspondingly, the V2X index spikes dramatically in those stress periods (peaking above 80% in late 2008 and in March 2020), reflecting surging uncertainty. Outside of crisis episodes, the V2X tends to range at more moderate levels (typically 10–30%), and it inversely co-moves with the SX5E: volatility rises when the stock market falls. This inverse relationship, visible in the figure, is characteristic of equity markets and underpins the motivation for using V2X as a forecast of risk, when the market sells off, implied volatility indices jump as investors brace for higher future volatility. The long time span plotted here underscores the variability in both series and the presence of distinct volatility regimes (low vs. high volatility periods), which any successful forecasting model must accommodate.

The chosen dataset provides a comprehensive view of both *realized* and *implied* volatility for a key equity market. Using the SX5E index prices allows us to compute realized volatility as the factual outcome against which forecasts can be benchmarked. Using the V2X index gives us the market's implied volatility forecast, which can be seen either as a competitor or complement to statistical models. This setup is well-grounded in the volatility forecasting literature. On one hand, past research finds that volatility is time-varying and forecastable to a significant extent, justifying the use of time-series models on historical prices. On the other hand, implied volatility incorporates forward-looking information and often outperforms traditional time-series forecasts, which is why we include the V2X measure. By studying both, we can assess whether model-based forecasts (such as those from GARCH-type models or other predictors introduced in later chapters) can match or improve upon the information in V2X. This data strategy aligns with many prior studies that evaluate the incremental value of volatility models relative to implied volatility indexes (e.g. comparing model forecasts of realized volatility to the VIX/VSTOXX).

2.1.4 Computing Daily Log Returns

For modelling volatility, we work with asset *returns* rather than raw prices. Returns are defined as the relative change in price, and we use the logarithmic (continuously compounded) return. Specifically, if P_t is the closing price (or index level) on day t, the log return is

$$r_t = \ln\left(\frac{P_t}{P_{t-1}}\right) = \ln(P_t) - \ln(P_{t-1})\,.$$

In practice, we compute this as the difference of the natural log of successive close prices for SX5E. The first day of each series has no return since we need two points to compute a return.

Log returns are used instead of simple returns (percentage change) for several reasons: (1) Log returns are time-additive over multiple periods (the k-day log return is just the sum of daily log returns), which simplifies multi-period analysis. (2) They make the distributional modelling easier. Many financial econometrics models (e.g. GARCH) assume returns have more tractable statistical properties when log-returns are used. In fact, Tsay (2010) notes that continuously compounded returns have "advantages over the simple net returns," particularly that multi-period log returns sum cleanly. (3) If prices follow a multiplicative random walk with small returns, log returns are approximately normally distributed (by the Central Limit Theorem over time), which suits many models. For small returns, $r_t \approx R_t$ (simple return), and modelling log-returns often yields stationary series with similar properties to price ratios.

In addition to SX5E returns, we can compute returns of V2X as $\ln(\text{V2X}t/\text{V2X}t-1)$ if needed (noting that V2X is an index of implied volatility, not a price). However, since V2X measures volatility itself, we may either use its level as a volatility proxy or use its returns in forecasting models. Either way, the same formula applies.

Thus, we compute two series of daily log returns:

$$r_t^{V2X} = \ln(V2X_t) - \ln(V2X_{t-1})\,.$$

Each is measured in natural log units. These log returns will be the inputs for volatility modelling (e.g. estimating GARCH models, calculating historical volatility, etc.). We will also present descriptive statistics (mean, variance, skewness, kurtosis) of these return series.

2.1.5 Data Cleaning: Missing Values and Outliers

Before analysis, the data were checked for any irregularities.

Weekends and public holidays (non-trading days) are simply missing from the dataset (no entry for those calendar dates). There is no need to artificially fill these: by using daily closing prices, we implicitly model on the set of actual trading days. We do not generate synthetic values for weekends. The time index of our data is simply the sequence of trading dates. We also inspect the return series for extreme outliers (extreme positive or negative returns). As expected, the largest moves correspond to crisis events (e.g. October 2008, March 2020). These are *real* market movements and not data errors, so we do not remove them. In fact, for volatility forecasting, these extreme observations are critical information. We do, however, verify that they do not reflect data entry errors (e.g. a misplaced decimal). The returns we computed were within plausible ranges: the minimum SX5E daily log-return was about –0.132 (–13.2%) and the maximum +0.104 (+10.4%), which are consistent with known crashes and rallies. Thus, no winsorization or outlier trimming is applied; GARCH models (if used) will account for heavy tails by allowing for fat-tailed error distributions if needed.

Since SX5E is an equity index, its constituents undergo corporate actions (dividends, splits, mergers) over time. However, the index methodology (free-float market cap weighting) automatically adjusts for these, and the published price index is typically on a price-return basis (dividends not included). Thus, no further adjustment for dividends is needed for analysing returns, because the index level already reflects splits and inclusion changes. It should be noted that index rebalances (twice yearly) do cause small jumps in level; technically these could introduce artificial shifts in returns. In practice, such rebalances cause tiny changes and are typically ignored in daily vol models. If one wanted to be thorough, a check for any unusually large jump on rebalance dates could be done, but for SX5E this is minor. Given the scope, we treat the index as exogenously given and do not perform adjustments beyond using the official closing values.

In summary, data quality is high. We have a nearly complete daily series for SX5E and V2X on trading days. Any missing high/low values are negligible, and the returns contain no suspicious outliers beyond market moves. Therefore, we proceed to analysis on this dataset as is, without elaborate cleaning. All further computations (log returns, vol measures) will be done on this cleaned series of closing prices.

2.1.6 Exploratory Data Analysis and Stylized Facts

We now examine the statistical properties of the two series (SX5E and V2X) to confirm expected stylized facts of financial time series. We focus on the log-return series for SX5E; similar observations often hold for V2X changes.

- *Descriptive statistics:* Table below summarizes key moments of daily log returns for SX5E and for V2X. For SX5E, the mean daily log-return is very close to zero (around 0.0001), reflecting that over short intervals prices fluctuate around no drift. The variance of SX5E log-returns is on the order of 1.8×10^{-4} (daily standard deviation ≈ 1.33%). This corresponds to an annualized volatility of roughly 21% ($1.33\% \times \sqrt{252}$). The distribution of SX5E returns is *leptokurtic*: the sample excess kurtosis is about 5.21 (raw kurtosis ≈ 8.21), far above 0 (normal kurtosis=3 has excess kurtosis=0). This indicates "fat tails" – extreme returns are more likely than a normal distribution would predict, a well-known stylized fact. The skewness is about –0.31, indicating a slight negative asymmetry (large negative returns are somewhat more extreme than large positives). These results are in line with empirical regularities: most stock index returns have excess kurtosis (Cont, 2001[53]) and often slight negative skew (losses tend to be sharp).

The V2X index (volatility index) has very different raw levels, but its log returns can be computed similarly. Its mean log-return is near zero, and its variance is larger (≈0.0044, daily SD ≈6.6%). The V2X returns are positively skewed (skew ≈ +0.78) and also leptokurtic (excess kurtosis ≈ 1.9). The positive skew indicates that increases in volatility (V2X jumps) are somewhat larger or more frequent than large downward moves. This asymmetry is sensible since volatility often jumps up in crises, but cannot go below zero; high volatility events can be very large while "calm" periods are bounded. Overall, V2X return distribution still shows fat tails.

These moment estimates empirically confirm two stylized facts: *heavy tails* (large kurtosis) and *asymmetric returns*. They justify the use of models (e.g. GARCH or other volatility models) that allow for non-normal and skewed shocks.

- *Time series plots and clustering:* A plot of daily SX5E log-returns over the sample (see Figure 2.2) reveals clear *volatility clustering*: periods of large fluctuations cluster together. For example, the autumn 2008 crisis period shows huge negative returns on consecutive days, and again March 2020. In contrast, mid-2017–mid-2019 had relatively small swings day-to-day. Such patterns are textbook volatile clustering: days of big returns (of either sign) tend to follow each other, producing runs of high volatility, while tranquil periods alternate. This phenomenon is well documented in financial data. It suggests that volatility is autocorrelated, indeed, the autocorrelation of absolute returns decays slowly (hyperbolically) rather than dropping to zero quickly. In other words, conditional on a turbulent day, near-term future days are more likely to be turbulent (and vice versa). This justifies using conditional heteroskedasticity models.

  The V2X series itself (index level) shows spikes during crisis. Plotting the V2X over time (not shown here) would reveal that it jumps in October 2008, spikes in late 2011 (Euro-debt fears), and in early 2020 (COVID). These are the same periods when SX5E volatility was high. The negative correlation between SX5E and V2X is apparent: when SX5E falls, V2X rises. This is the well-known *leverage effect* or "flight-to-safety" dynamic: equity crashes tend to coincide with fear rising (volatility index surging). Quantitatively, during 2005–2024

[53] Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), pp. 223–236.

the contemporaneous correlation between SX5E daily return and V2X change is strongly negative (empirically around –0.5 to –0.7). This negative link is noted in the literature: volatility indices like VSTOXX/V2X move *up* when equity markets decline. The economic rationale (often called the "leverage effect") is that falling equity values increase financial leverage and risk, or simply that bad news raises uncertainty more than good news lowers it.

- *Stylized facts in detail:* Summarizing, our data exhibit the "stylized facts" common in equity returns: (1) *Lack of linear autocorrelation* in raw returns: the daily log returns themselves have almost zero autocorrelation (efficient market property); this implies predictability of future returns is very limited. (2) *Volatility clustering:* strong serial correlation in the magnitude (squared or absolute) of returns. This is visible in the time series and confirmed by computing autocorrelations of squared returns (significant persistence). (3) Heavy tails: returns have high kurtosis, meaning extreme events are more probable than under normality. (4) Leverage (asymmetry): negative returns tend to precede higher conditional volatility, consistent with our negative skew and the index correlation. (5) Volatility risk premium: As noted in regulatory reports, implied volatility (V2X) on average exceeds realized historical volatility, indicating a risk premium investors demand for volatility. In our sample, the historical 30-day volatility of SX5E (rolling standard deviation) is typically lower than the V2X level. This gap is itself informative about market sentiment.

  These empirical features justify standard volatility modelling techniques (such as GARCH, EGARCH, or stochastic volatility models) that accommodate clustering and non-normality. We will build on these facts by using GARCH-family models to forecast volatility. The descriptive statistics here confirm that both SX5E and V2X have rich dynamics and that volatilities are time-varying and predictable to some extent.

- *Correlation and co-movement:* Briefly, we note that SX5E and V2X are strongly linked. A simple scatter plot of SX5E returns against V2X changes (not shown) reveals the negative relationship. During stable periods (low V2X), returns are small on average; when V2X jumps above ~20-30, negative returns on the equity index become likely. This supports using V2X as a covariate or additional predictor in volatility forecasting. In fact, we might later estimate models of volatility that include the contemporaneous or lagged V2X level, treating it as a forward-looking volatility signal (implied vol) versus realized vol.

The exploratory analysis confirms that our datasets exhibit the expected properties of stock index returns in a major market: unpredictable mean returns, volatility clustering, fat tails, and leverage effects. This validates the choice of log-return transformation and sets the stage for formal volatility modelling.

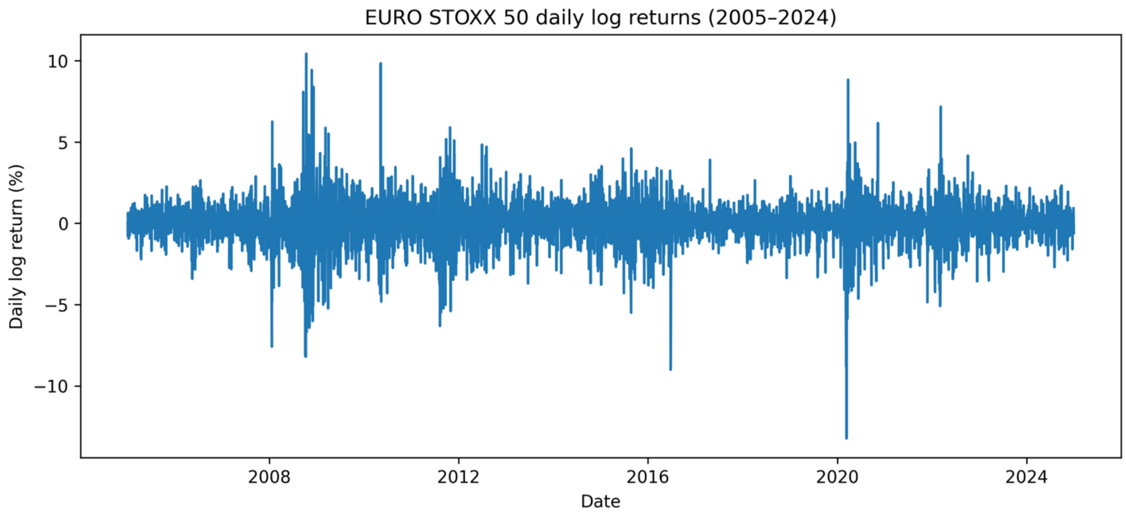

Figure 2.2 plots the daily log returns of the EURO STOXX 50 Index over the 2005–2024 period. Several features stand out. Large negative outliers occur during the 2008 crisis (Oct 2008) and the COVID crash (Mar 2020), confirming that these events generated the largest shocks to the market. Between crises, the returns fluctuate in a relatively narrow band. This temporal clustering of volatility can be seen by the "fatter" appearance of the series in crisis intervals versus calm periods. A corresponding plot of the V2X series show spikes peaking in those same crisis times, mirroring the inverse of the return spikes. Such visualizations underscore the negative correlation and the episodic nature of volatility.

2.1.7 Choice of Indices Over Alternatives

Finally, we justify why the EURO STOXX 50 and V2X indices are more appropriate for our study than alternative measures. Using a single firm's stock (or an average of firms) would not capture aggregate market volatility; IPO outcome studies generally focus on market-wide conditions, not firm-specific risk. The SX5E index aggregates 50 major firms, providing a broad market representation. Volatility forecast from SX5E reflects systemic Eurozone risk that can influence all firms. An individual stock's volatility might be driven by idiosyncratic news, which is not relevant for general market predictions. Thus, using a market index avoids idiosyncratic noise and ensures relevance to average IPO issuer conditions.

One could consider, say, Germany's DAX or France's CAC 40 instead. While these capture large economies, the EURO STOXX 50 spans the entire Eurozone (major economies including Germany, France, Italy, Spain, etc.). IPO activity in Europe is not limited to one country. By using the pan-Euro index, we capture volatility common to the region. Indeed, regulation and financial linkages mean shocks tend to spread across borders in the Eurozone. The SX5E is also more liquid and has more futures/options volume than any single-country index, making V2X more robust.

Also, the STOXX Europe 600 includes non-Euro EU markets (UK, Switzerland, etc.) and smaller firms. We prefer Eurozone-specific volatility, so SX5E is more focused on the target geography. Moreover, SX5E constituents are blue-chip, ensuring the index is not overly influenced by small-cap noise.

Then, sector indices (like STOXX Banks or Tech) are too narrow and volatile for forecasting broad IPO trends, which are driven by overall market sentiment. By using the overall index, we capture aggregate risk. Sector-specific volatility could be considered in future research but is outside our scope.

Additionally, the only direct alternative to V2X would be another volatility measure, such as the VDAX (Germany), or a realized volatility measure from historical data. We specifically use V2X because it is the canonical forward-looking implied vol for the Euro STOXX 50. It is traded on Eurex and is recognized by regulators (e.g. ESMA cites VSTOXX as the EU's volatility benchmark). It reflects option market expectations, which often contain predictive information. While one could compute realized volatility (e.g. historical variance) from the SX5E returns, using the implied volatility index can capture sentiment not yet realized. That is, V2X is forward-looking. For forecasting IPO outcomes, forward-looking vol may contain additional signals beyond what one would get from past returns alone.

In sum, the combination of the EURO STOXX 50 and its volatility index (V2X/VSTOXX) offers a comprehensive picture of Eurozone market conditions. This choice is supported by literature: for instance, Cont (2001) and others emphasize studying index-level returns for broad market analysis, and regulatory sources identify VSTOXX as the standard European volatility benchmark. The indices we use have the desirable property of negative correlation and consistent stylized behaviour, making them well-suited for modelling volatility over the chosen sample.

Series	Mean (×10^-4)	Variance (×10^-4)	Skewness	Kurtosis	Obs.
SX5E log-ret	0.97	1.78	–0.31	8.21	5092
V2X log-ret	0.32	43.74	+0.77	4.01	5092

Table 2.1 Descriptive statistics of SX5E and V2X returns

2.2 Historical Volatility Models

Volatility is not directly observable and must be proxied using observable price data. A natural starting point is historical volatility, defined as a function of past returns. Historical measures are widely used by practitioners because they are transparent, easy to compute, and require only time-series data on prices. In this thesis, historical volatility estimates serve two roles: (i) they provide simple baseline forecasts of market uncertainty, and (ii) they act as benchmarks against which we compare forward-looking measures such as implied volatility and parametric forecasts from GARCH-family models (introduced later).

Let P_t denote the EURO STOXX 50 closing level (SX5E) on trading day t, and let daily log returns be:

$$r_t = \ln(P_t) - \ln(P_{t-1})$$

Throughout this section, volatility estimates are expressed as *annualized percentages*, using 252 trading days per year:

$$\sigma^{\text{ann}} = \sigma^{\text{daily}\sqrt{252}} \times 100$$

In addition, unless otherwise stated, a "30-day" window refers to 30 trading days (we later discuss horizon alignment with option-implied 30-calendar-day measures).

2.2.1 Rolling-Window Volatility

A standard estimator of historical volatility is the rolling standard deviation of returns computed over a fixed window of the most recent H observations. For a window length H, the rolling variance estimator at time $t$is:

$$\widehat{\sigma_{t,H}^2} = \frac{1}{H-1}\sum_{i=0}^{H-1}\left(r_{t-i} - \overline{r}_{t,H}\right)^2, \quad \overline{r}_{t,H} = \frac{1}{H}\sum_{i=0}^{H-1} r_{t-i}$$

and the corresponding rolling volatility estimate is

$$\widehat{\sigma_{t,H}} = \sqrt{\widehat{\sigma_{t,H}^2}}\ .$$

Because daily index returns have very small means over short horizons, results are typically very similar if one sets

$$\overline{r_{t,H}} \approx 0\ .$$

In the empirical analysis, we compute rolling volatility using $H = 30$ trading days. This choice provides a smooth measure of "recent" volatility and is commonly used as a practical monthly-type window. In our 2005–2024 sample, the mean 30-day rolling volatility of SX5E returns is approximately 1.19% per day, corresponding to roughly 18.8% annualized volatility.
Rolling-window volatility has a simple forecasting interpretation: it can be used as a naïve one-step-ahead forecast,

$$\hat{\sigma}_{(t+1|t)} \equiv \hat{\sigma}_{t,H}$$

i.e., tomorrow's volatility is forecast as today's rolling estimate. Its key limitation is that it reacts mechanically and with delay to sudden regime changes: large shocks only affect the estimator gradually as they enter (and later leave) the rolling window.

2.2.2 Exponentially Weighted Moving Average

A closely related historical approach is the Exponentially Weighted Moving Average (EWMA) model, popularised in practice by RiskMetrics. EWMA assigns larger weights to recent observations and exponentially decaying weights to older observations, allowing volatility estimates to adapt more quickly to changing conditions than a simple rolling window.

EWMA is typically specified in variance form:

$$\hat{\sigma}_t^2 = \lambda \hat{\sigma}_{(t-1)}^2 + (1-\lambda) r_{(t-1)}^2,$$

$$0 < \lambda < 1,$$

where λ controls the persistence (decay) of volatility. A common choice in the literature and in risk management is $\lambda = 0.94$ for daily data. Intuitively, this recursion says that today's conditional variance is a weighted average of yesterday's variance and yesterday's squared return. With $\lambda = 0.94$, shocks decay gradually (the half-life of a squared-return shock is about 11 trading days), producing a responsive but still smooth volatility series.

Using $\lambda = 0.94$, the mean EWMA volatility for SX5E returns over 2005–2024 is approximately 1.19% per day (about 18.9% annualized), very close to the rolling-window average, but with faster adjustment around turning points. Figure 2.3 illustrates this visually: EWMA typically rises more rapidly at the onset of stress episodes and falls more quickly in the early stages of recovery, while the rolling estimator can lag due to its equal-weight window.

EWMA also admits a forecasting interpretation similar to rolling volatility:

$$\hat{\sigma}_{(t+1|t)} \equiv \hat{\sigma}_t \,,$$

making it a convenient benchmark forecast in out-of-sample comparisons.

In our sample (2005–2024), the rolling 30-day volatility and EWMA volatility series are very similar in overall level but differ slightly in responsiveness. The rolling 30-day annualized volatility has a mean of 18.74% and a median of 15.89%, with a 95th percentile of 38.55%. The EWMA annualized volatility ($\lambda = 0.94$) has a mean of 18.82% and a median of 16.27%, with a 95th percentile of 38.27%. Both estimators capture the same volatility regimes and crisis spikes, while EWMA tends to adjust marginally faster at turning points because it assigns greater weight to recent squared returns.

2.2.3 Range-based volatility estimators

In addition to close-to-close returns, the SX5E dataset contains daily highs and lows, which can be used to construct range-based volatility proxies. A classic example is the Parkinson (1980) estimator, which exploits the intraday high–low range:

$$\widehat{\sigma^2_{\text{Park},t}} = \frac{1}{4\ln(2)}\left[\ln\left(\frac{H_t}{L_t}\right)\right]^2,$$

where H_t and L_t denote the daily high and low index levels. Range-based estimators can be more statistically efficient under idealized assumptions (continuous diffusion without jumps), but they may be sensitive to market microstructure effects and jumps. For this reason, in the baseline analysis we rely on close-to-close return-based measures (rolling and EWMA), while range-based measures can be used as *robustness checks* where appropriate.

```
Rolling 30d vol (ann, %)  mean/median/p95: 18.73960578601015 15.8932278681419 38.55023920712116
EWMA vol (ann, %)         mean/median/p95: 18.822712765545706 16.26639064084733 38.266329240296926
```

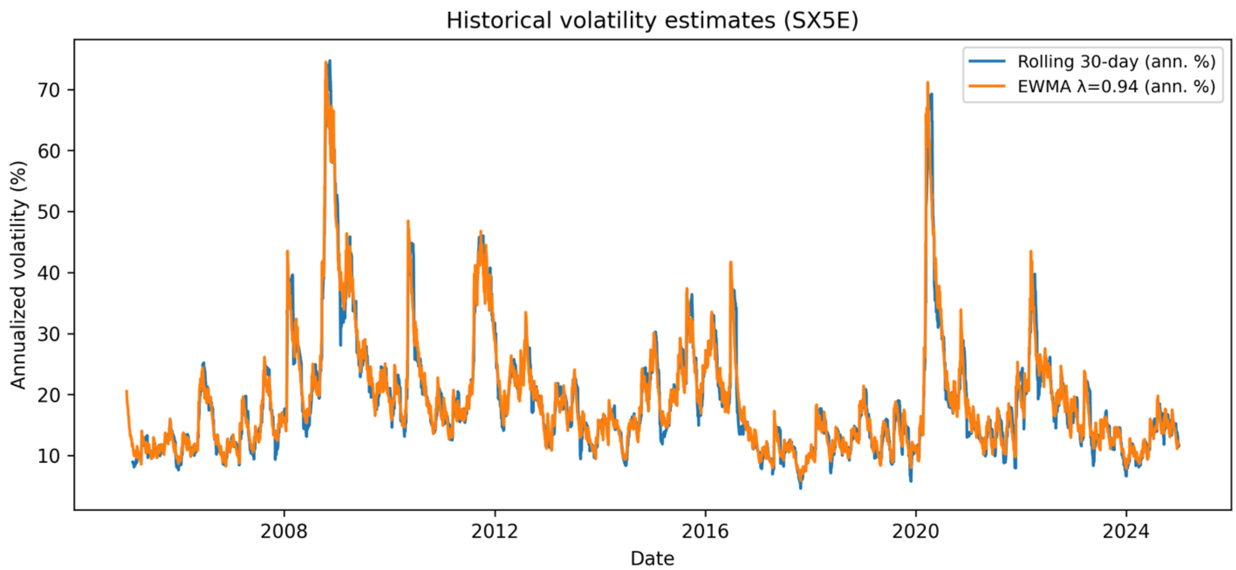

Figure 2.3 compares rolling 30-day volatility (equal-weighted window) with EWMA volatility (λ = 0.94). Both measures capture volatility clustering and crisis spikes, while EWMA reacts slightly faster to new shocks.

Overall, historical volatility estimators provide a simple and intuitive quantification of market uncertainty, and they capture well-known features such as volatility clustering and regime shifts: both rolling and EWMA volatility rise sharply during crisis periods (e.g., 2008-2009 and 2020) and revert toward lower levels in calmer periods. However, these estimators are entirely *backward-looking* and do not incorporate information from derivatives markets or explicitly model conditional variance dynamics.

For this reason, historical measures here are treated as baseline tools. In later sections, we compare their forecast performance against *implied volatility* from the options market and against *GARCH-family conditional variance models*, and we evaluate which measures provide the most informative real-time representation of the financing environment faced by IPO issuers.

2.3 Implied Volatility and the V2X Index

Implied volatility represents the market's consensus expectation of future volatility, derived from options prices on an underlying asset. In essence, it is the volatility figure that, when plugged into an option pricing model (such as Black-Scholes), yields the observed market price of the option. As a forward-looking measure, implied volatility is often interpreted as the market's "best guess" of upcoming volatility over the option's remaining life. Because implied volatility tends to spike during market downturns (when options prices surge due to demand for protection), it has earned the moniker of an "investor fear gauge". In his seminal work, Whaley (2000)[54] introduced the CBOE Volatility Index (VIX) as a measure of implied volatility on the S&P 500 and dubbed it the "Investor Fear Gauge," noting that "expected stock market volatility rises when the market declines," thus linking high index values to heightened investor anxiety. In other words, implied volatility indices like the VIX are widely viewed as barometers of market sentiment and stress, rising in times of uncertainty and declining during stable periods. Implied volatility is also commonly treated as a forecast (albeit a biased one) of actual future volatility. Carr and Wu (2006)[55] note that markets regard implied volatility indices as both a forward-looking forecast of subsequent realized volatility and an indicator of market stress. This dual interpretation, as a predictor of volatility under normal conditions and as a gauge of fear during turmoils, underpins the extensive use of implied volatility in both academic research and industry practice.

2.3.1 V2X as a Model-Free Implied Volatility Index

We already mentioned how V2X, is the Eurozone's equivalent of the VIX and is constructed in an analogous manner. Importantly, the V2X is a "model-free" implied volatility index, meaning it is not produced by inverting a single option pricing model for at-the-money volatility, but rather by aggregating information from the entire spectrum of option prices across strikes. Specifically, the index is calculated by taking a weighted portfolio of out-of-the-money EURO STOXX 50 put and call options, spanning a range of strike prices, such that this portfolio's price is proportional to the risk-neutral expected variance over the next 30 days. Mathematically, the V2X (like the VIX) is derived from the option price integrals that replicate a forward-start variance swap, and the index value quoted is the annualized standard deviation (volatility) corresponding to that risk-neutral 30-day variance. Taking the square root of the fair variance swap rate gives the implied volatility index level. This model-free construction was outlined by Carr and Wu (2006) in their description of the

[54] Whaley, R.E. (2000) 'The investor fear gauge', *Journal of Portfolio Management*, 26(3), pp. 12–17.
[55] Carr, P., & Wu, L. (2006). A tale of two indices. *Journal of Derivatives, 13*(3), 13–29.

“new VIX” which uses all available strikes and a static replication formula to capture variance. The benefit of this approach is that it does not assume any specific model (e.g. Black–Scholes) for the underlying’s dynamics, and hence the index truly reflects market prices (and expectations) of volatility directly from option quotes, without model error. The V2X, calculated by STOXX Ltd. and disseminated in real time, thus represents the risk-neutral expected volatility of Eurozone blue-chip stocks over the next month.

As a volatility index, the V2X encapsulates two key pieces of information: (1) the market’s expectation of future realized volatility of the EURO STOXX 50 (often called the “physical” or actual expected volatility under the real-world measure), and (2) a risk premium for volatility (since option prices are set under the risk-neutral measure). Because investors are typically risk-averse and pay a premium for options as insurance, the V2X is generally higher than the actual volatility that subsequently realizes, reflecting this variance risk premium component. In other words, the V2X can be viewed as the sum of expected stock market uncertainty plus an extra compensation for bearing variance risk. We will discuss this volatility risk premium in detail later in this section. First, we examine the behaviour of the V2X time series and its informational content.

2.3.2 Time Series Behaviour of the V2X Index

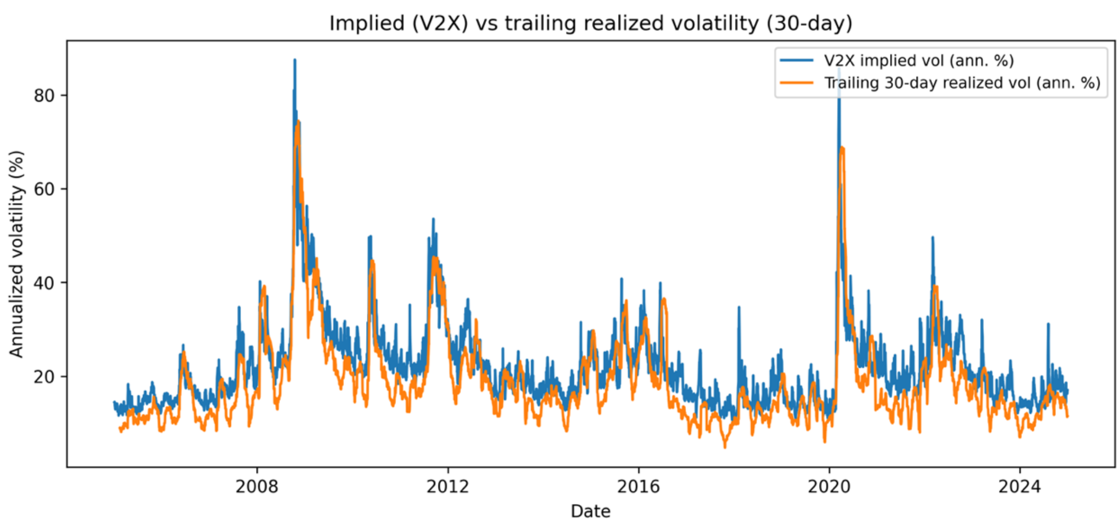

Figure 2.4 compares implied volatility with 30-day trailing realized volatility over the 2005–2024 period

Time series of the EURO STOXX 50 Volatility Index (V2X) from January 2005 to December 2024. Major spikes align with periods of market turmoil: the 2008 Global Financial Crisis, the Eurozone sovereign debt crisis (2011), and the COVID-19 crash (March 2020). The index is mean-reverting, with elevated volatility during crises reverting back toward average levels over subsequent months.

Over the 2005-2024 period, the V2X index exhibited a highly variable yet mean-reverting trajectory, punctuated by sharp spikes during major financial crises. The long-run average level of V2X over this period is around the mid-20s (approximately 24% in annualized volatility terms). In stable market conditions, the index has often hovered in the teens or low-20s, reflecting modest anticipated volatility. For instance, in early January 2020, just before the COVID-19 shock, the V2X reached as low as 10.7, one of its lowest readings on record. Such low values indicate markets pricing in very low near-term volatility (often a sign of complacency or stable outlooks)

However, during periods of acute market stress, the V2X has soared to extreme levels, reflecting surging demand for options and drastically higher expected volatility. The Global Financial Crisis (GFC) of 2008 saw the V2X reach an all-time high: in October 2008, as equity markets plunged amid panic, the V2X spiked above 85 (its highest level in the historical sample). This value signifies an expected volatility of over 85% per annum, an extraordinary level consistent with the turmoil at that time. Another significant episode was the Eurozone sovereign debt crisis around 2010–2012. The V2X climbed to about 53 in May 2010 (during the initial Greek debt panic), and it surged again in late 2011, peaking around 49-50 in September 2011 when fears of a Eurozone breakup intensified. These Eurozone-focused crises produced volatility spikes nearly double the index's long-run mean, though not as extreme as 2008. More recently, the COVID-19 market crash in March 2020 triggered another record-breaking volatility spike. Within a matter of days, the V2X exploded from historically low levels to around 85.6 on 16 March 2020, nearly matching the 2008 peak. This was the highest level since the GFC, as the pandemic's sudden economic shock led to "unprecedented panic" in options markets. The speed of this spike was remarkable, on 12 March 2020 alone, the V2X jumped by the largest one-day amount ever recorded, coinciding with a -12.4% collapse in the EURO STOXX 50 index (its worst day on record).

Each volatility spike has been followed by a reversion toward more normal levels, underscoring the mean-reverting nature of volatility. For example, after the October 2008 peak, the V2X gradually subsided over the next year: it took about 13 months for the index to fall back to its long-run average following the GFC's peak. In the aftermath of the Eurozone crisis spike of late 2011, the V2X also retreated to average levels in roughly 4 months, aided by central bank interventions that calmed investors. Following the COVID-19 volatility explosion in March 2020, the index likewise began to decline once extraordinary policy responses were enacted, though it remained elevated for some time. These patterns align with a well-known characteristic of volatility: periods of extreme volatility tend not to last indefinitely; volatility "reverts to the mean" as crises abate and markets normalize. Indeed, volatility's tendency to cluster in high-vol regimes and then subside is evident from the V2X time series. Between the crisis episodes, there are extended stretches of lower volatility, but once a shock hits, volatility can jump sharply and then gradually ease off rather than staying high permanently.

Overall, the time series of V2X reflects the episodic nature of market uncertainty. It remains relatively low and range-bound in tranquil times, and it experiences short, intense upward bursts during market upheavals. These dynamics justify the view of implied volatility indices as "fear gauges": they essentially trace the ebbs and flows of investor fear. In calm periods, fear is low (volatility is low); in crises, fear is elevated (volatility spikes). We next examine the statistical relationship between the V2X index and the equity market, to quantify this fear gauge intuition.

2.3.3 V2X as a Fear Index: Correlation with Equity Market Returns

A striking empirical regularity observed in our data is the strong inverse relationship between implied volatility and equity index returns. Daily changes in the V2X index are highly negatively correlated with returns on the EURO STOXX 50 index. In our sample from 2005-2024, the contemporaneous correlation between V2X and SX5E daily returns is about -0.80, indicating that when the stock market falls, implied volatility almost invariably jumps, and vice versa (a manifestation of the "leverage effect"). This nearly -80% correlation is in line with prior findings for U.S. markets, where the VIX and S&P 500 returns also exhibit correlations around -0.7 to -0.8. The economic intuition is straightforward: a market decline (especially a sharp one) prompts investors to scramble for protection (buying put options or volatility products), which drives up option-implied volatility. Moreover, due to leverage effects and volatility feedback, a drop in equity prices can mechanically increase future return volatility, reinforcing the inverse relation. Carr and Wu (2006) confirm this strong contemporaneous negative correlation between VIX and SPX returns, attributing it to the classic leverage effect first noted by Black (1976)[56]. Thus, the "fear index" aspect of V2X is quantitatively supported: on days when stocks plunge, V2X typically spikes (fear rises), and on days when stocks rally, V2X subsides (fear recedes).

```
                     count       mean       std        min        25%  \
V2X                 5062.0  22.017774  8.772760  10.678300  16.026900
Realised30_past_pct  5062.0  18.828141  9.912153   4.588219  12.599484
Realised30_fwd_pct   5062.0  18.828663  9.911703   4.588219  12.599484
VRP_past_pct         5062.0   3.189633  4.961999 -27.973362   1.079469
VRP_fwd_pct          5062.0   3.189111  4.811096 -28.411194   1.091345

                          50%        75%        max
V2X                 20.058500  25.219400  87.512700
Realised30_past_pct  16.005186  21.815693  74.733151
Realised30_fwd_pct   16.005186  21.815693  74.733151
VRP_past_pct          3.543539   5.796721  33.537971
VRP_fwd_pct           3.537721   5.761856  32.703756

Correlations (levels):
Corr(V2X, Realised30_past_pct) = 0.8659
Corr(V2X, Realised30_fwd_pct)  = 0.8744

Daily relationships:
Corr(ΔV2X, r)     = -0.7976
Corr(ΔV2X, |r|)   = 0.1968

Lead-lag (yesterday V2X vs today move):
Corr(V2X_t-1, r_t)   = 0.0312
Corr(V2X_t-1, |r_t|) = 0.4809
```

Figure 2.5: Summary statistics and key correlations between V2X implied volatility, realized volatility measures, and the volatility risk premium (Full sample 2005–2024). The upper panel reports distributional moments for each series. The middle panel documents the contemporaneous

[56] Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, pp. 177–181.

daily relationship between V2X changes and SX5E returns: the correlation of -0.80 confirms the strong inverse co-movement characteristic of equity implied volatility indices. The lower panel examines lead-lag dynamics: the near-zero correlation between yesterday's V2X and today's return direction (0.03) contrasts with the moderate correlation with today's absolute return magnitude (0.48), indicating that implied volatility predicts the size of future market moves but not their sign.

It is important to note, however, that this relationship is largely contemporaneous. If we lead or lag one series relative to the other, the correlation largely dissipates. For instance, today's V2X level has virtually no correlation (~0.03 in our data) with the next day's stock return. Likewise, yesterday's stock return has only minimal predictive power for today's volatility index movement. Empirically, we find a tiny positive lagged correlation (~+0.03) between V2X and subsequent SX5E returns, which is statistically negligible. This implies that implied volatility changes do not *cause* or predict equity price moves on a one-day-ahead basis, nor vice versa – instead, both tend to react simultaneously to common news or shocks. Carr and Wu (2006) report that while lagged stock returns have marginal predictive power for next-day VIX movements, lagged VIX changes do not predict future stock returns. In other words, volatility jumps and price drops occur together, consistent with volatility being largely an endogenous response to market movements (or both responding to a third factor, such as new information). The lead-lag effects are minimal: volatility indices adjust so quickly to news that any information content for subsequent equity returns is quickly impounded.

While implied volatility itself does not predict the *direction* of market returns, it does contain information about the magnitude of future market moves. We find that the V2X is moderately positively correlated (~0.48) with the *next day's absolute* equity return (i.e. with the magnitude of returns, regardless of sign). In practical terms, if the V2X is elevated today, tomorrow's trading range is likely to be wider than usual (though the price could move up or down). This aligns with the idea that implied volatility foretells future market volatility. Indeed, options-based volatility indices are often interpreted as forecasts of variance. In our analysis, days of high V2X tend to be followed by days of high realized volatility (large absolute returns), indicating that V2X contains a forward-looking element regarding near-term turbulence. This is consistent with broader evidence that implied volatility can predict future realized volatility to a significant degree. Later in this section, we delve deeper into implied volatility's forecasting ability relative to historical volatility measures. But already, the strong correlation between V2X and realized volatility measures (~0.87 in our sample) suggests that V2X moves hand-in-hand with the actual volatility being experienced (or soon to be experienced) in the market. The STOXX research on VSTOXX likewise notes that volatility indices tend to co-move with realized market swings, exhibiting well-known features such as negative correlation with the underlying index, volatility clustering, and mean reversion.

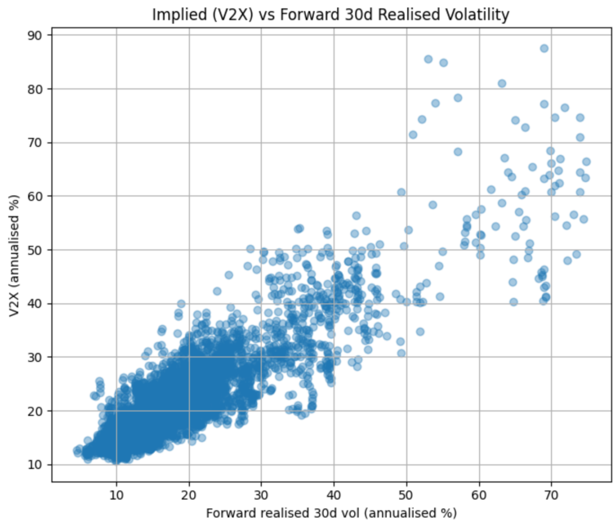

Figure 2.6: Scatter plot of *implied volatility (V2X) vs 30-day forward realised volatility (annualised %).*

The V2X index empirically earns its reputation as Europe's "fear gauge". Its level is inversely related to equity market performance (fear up when markets fall), and while it does not predict market direction, it does provide information on the anticipated magnitude of market fluctuations. These properties make implied volatility an invaluable indicator for market participants. Next, we introduce the concept of the volatility risk premium, which is intricately linked to implied volatility's behaviour and is key to understanding why implied volatility tends to *exceed* realized volatility on average.

2.3.4 The Volatility Risk Premium (VRP)

An important insight from comparing implied and realized volatilities is that implied volatility is usually higher than subsequent realized volatility. Investors are, on average, willing to pay a premium for volatility insurance, meaning they accept paying option prices that imply a higher volatility than they actually expect to occur. This difference is known as the Volatility Risk Premium (VRP). We define the VRP here as V2X minus realized volatility (for a given horizon), which can be thought of as the excess implied volatility above the actual volatility realized. In our context, we computed the VRP using both past 30-day realized volatility and forward 30-day realized volatility for the EURO STOXX 50.

In both cases, the V2X tends to exceed the realized metric, yielding a generally positive VRP.

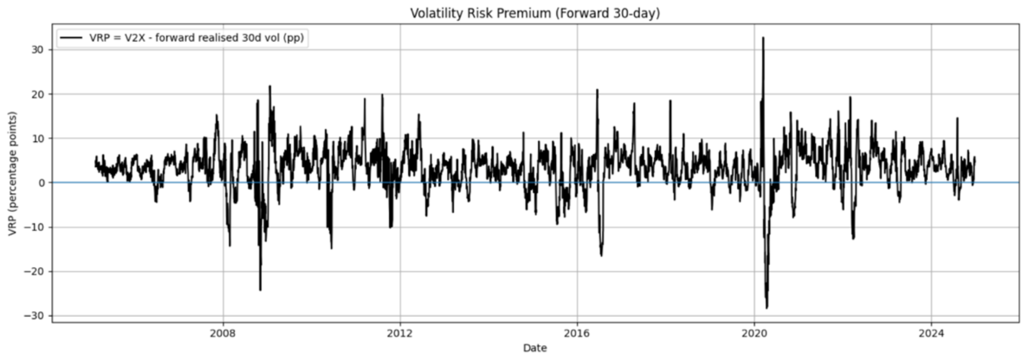

Figure 2.7: *VRP forward time series*

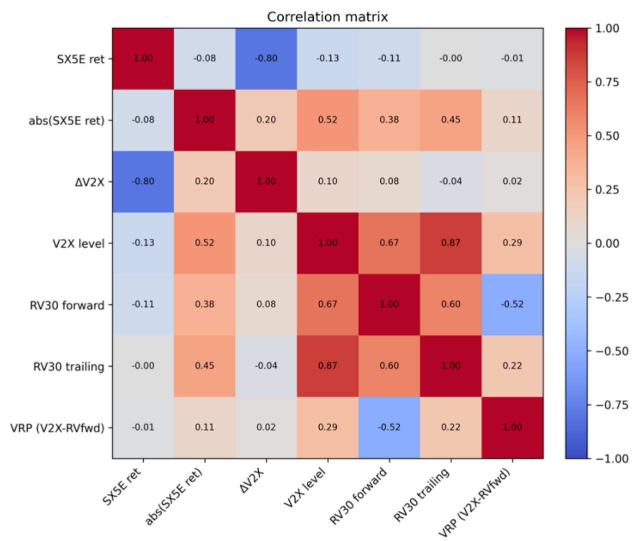

Figure 2.8 presents the correlation matrix between implied volatility, trailing and forward realized volatility, historical volatility estimators (Rolling(30), EWMA(0.94)), and 30-day GARCH-family forecasts. Historical measures are nearly perfectly correlated with one another, reflecting their shared backward-looking structure. Implied volatility is strongly correlated with trailing realized volatility but exhibits more moderate correlation with forward realized volatility, highlighting the inherent uncertainty in forecasting future volatility. GARCH-family forecasts are highly correlated among themselves and closely track persistent volatility regimes.

The existence of a persistent positive VRP (implied > realized) reflects the fact that volatility under the risk-neutral measure (implied by option prices) is higher than volatility under the physical (real-world) measure. This is a direct consequence of risk-averse investors requiring compensation to bear variance risk. In equilibrium, option sellers (who are short volatility exposure) demand a premium, which manifests as higher option prices (hence higher implied volatilities) relative to the actual expected volatility. Extensive empirical research has documented this phenomenon. For example, Bakshi et al. (2003)[57] find strong evidence of a negative volatility risk premium in equity index options, meaning that a strategy of selling options (delta-hedged) yields positive average profits, equivalently, the buyer of volatility pays an insurance premium. In their study, delta-hedged option portfolios significantly underperform zero (negative average returns), especially during high-volatility periods, which is consistent with implied volatilities being set above subsequent realized volatilities. Likewise, Carr & Wu (2006) estimate that the variance risk premium for 30-day index variance is highly significantly negative, on the order of 2000+ basis points in variance terms (implied variance exceeding realized variance by about 21.6% on average in annualized variance). When translated to volatility terms, this corresponds to implied volatility (VIX) being roughly 2 volatility percentage points higher on average than actual 30-day volatility of the S&P 500. They interpret this highly negative variance risk premium as evidence that investors are willing to pay a substantial premium for insurance against volatility spikes. In our Euro Stoxx 50 data, we similarly observe that V2X exceeds realized volatility by a notable margin on average, implying a persistent volatility risk premium in the Eurozone options market.

The economic meaning of the volatility risk premium is tied to risk aversion and the distribution of returns. Volatility (variance) is not directly tradable, but through options and variance swaps one can take positions on future volatility. A negative VRP (implied > realized) means that being short volatility (e.g. selling variance swaps or options) has yielded positive returns on average, compensating those sellers for the risk they bear. The flip side is that investors buying protection (long volatility) are willing to lose money on average essentially paying an insurance premium to transfer risk. The premium is high partly because volatility shocks are correlated with bad times (market crashes), and also because volatility payoff profiles are negatively skewed (shorting volatility yields small steady gains punctuated by large losses in crises). Thus, a high average return (premium) is required to induce investors to provide volatility insurance. This aligns with the notion that "investors are averse to variations in return variance", and therefore the mean volatility under the physical distribution is lower than under the risk-neutral distribution. In practical terms, the VRP can be seen as a proxy for market risk aversion and tail-risk compensation. It tends to widen in times of stress (implied volatility shoots up far more than actual volatility, as fear and demand for insurance peak) and can narrow or even invert in exceptionally calm or optimistic conditions (though a consistently negative VRP – realized > implied – is rare and typically short-lived).

From a statistical perspective, the VRP is often studied as variance risk premium = implied variance – expected realized variance. Our use of both past and forward realized volatility in computing VRP provides two views: one ex-post (with forward realized) and one ex-ante (with trailing realized as a proxy for expected volatility). Both confirm that V2X tends to price in more volatility than has been observed or even anticipated based on recent history. This has important

[57] Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies, 16*(1), 101–143.

implications: the size of the VRP has been shown to have predictive power for future returns. A large body of literature (e.g., Bollerslev, Tauchen & Zhou 2009[58]; Bekaert, Hoerova & Lo Duca 2013[59]) finds that a high variance risk premium is a predictor of higher future stock returns, consistent with a compensation for risk. Intuitively, when implied volatility is *much* higher than current realized volatility, indicating investors are very fearful and willing to pay up for protection, subsequent stock returns tend to be elevated (possibly as the fears mean-revert or as risk premia are earned). We will not delve deeply into return predictability, but it is worth noting that the VRP is a key state variable in asset pricing, linking option-implied information to the equity risk premium. By contrast, the level of volatility itself (without adjusting for the risk premium) has little predictive power for equity market returns. This highlights that it is *excess* volatility pricing (fear) rather than volatility per se that matters for expected returns.

In summary, the volatility risk premium reflects the systematic gap between implied and realized volatility. Its consistently positive magnitude (for implied minus realized) underscores that implied volatility indices like V2X are not pure forecasts of actual volatility but are risk-adjusted forecasts. For the purposes of this thesis, understanding the VRP is crucial: it reminds us that implied volatility is an upward-biased predictor of real volatility, with the bias representing compensation for risk. In subsequent sections, we will account for this when comparing implied volatility to other forecasting methods.

2.3.5 Implied vs. Historical Volatility: Forward-Looking Information Content

A central question in volatility analysis is how implied volatility (like V2X) compares to historical volatility estimators in terms of forecasting future volatility. Historical measures might include realized volatility (e.g. the standard deviation of past 30-day returns) or model-based forecasts (such as GARCH or other time-series models). Implied volatility has a key conceptual advantage: it aggregates the wisdom of market participants about future volatility, incorporating new information in real time. In contrast, historical estimates rely solely on past price data. Thus, implied volatility is inherently forward-looking, whereas realized volatility (by definition) is backward-looking. The question is whether this forward-looking metric indeed contains superior or complementary information about the future.

Empirical studies have largely found that implied volatility is an informative, though biased, predictor of future realized volatility. In other words, while implied vol tends to overestimate actual volatility (due to the VRP bias), it often carries more information about future movements than purely historical measures. For example, Christensen and Prabhala (1998)[60] showed early on that implied volatility subsumes much of the information in past volatility when predicting future

[58] Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies, 22*(11), 4463–4492.
[59] Bekaert, G., Hoerova, M., & Lo Duca, M. (2013). Risk, uncertainty and monetary policy. *Journal of Monetary Economics, 60*(7), 771–788.
[60] Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics*

volatility (for individual stocks and indices). More directly related to implied indices, Carr & Wu (2006)[61] report that the VIX index has significant predictive power for subsequent realized variance of the S&P 500. They find that once VIX is included as a regressor in a volatility forecasting model, GARCH-type volatility estimates "do not provide extra information". This suggests that the options market's implied volatility already encapsulates the relevant volatility clustering, mean-reversion, and leverage effect information that econometric models attempt to capture. Our own results are consistent with this view: the V2X's strong correlation with forward 30-day realized volatility (about 0.70 in our data) indicates that it tracks future volatility movements reasonably well, certainly better than a naive historical average would. Moreover, implied volatility reacts faster to market events than realized volatility does. When a sudden shock occurs, V2X will jump immediately (even before the underlying index has realized all the volatility, since options price in expected future turmoil). Realized volatility, often computed over a rolling window, will lag the jump, rising only as the large returns filter into the window. Thus, V2X often leads realized volatility at turning points, a quality that is valuable for forecasting. Indeed, we saw evidence of this in the lead-lag correlations: V2X has a moderate correlation with next-day absolute returns (~0.48), whereas yesterday's realized volatility would have a much weaker correlation with tomorrow's volatility. This implies implied vol contains a forward-looking element that pure historical vol lacks.

That said, implied volatility is not a perfect predictor of realized volatility. As noted, it is generally biased upward (owing to the risk premium). Additionally, if there is transient option market demand (e.g., hedging flows) or measurement noise, implied vol can sometimes overshoot or undershoot actual volatility that materializes. Nonetheless, many studies have concluded that implied volatility is typically the single-best predictor of future volatility available, in a root-mean-squared error sense, even outperforming sophisticated econometric models – especially over short horizons like 1-month. For example, research by Poon and Granger (2003)[62] and others finds that implied volatility contains unique information not contained in past returns and thus improves volatility forecasts. This is intuitive: options traders collectively process all available public (and perhaps some private) information about future risks (including macroeconomic outlooks, earnings, political events) which may not be reflected in the historical price series yet. Implied volatility indexes therefore serve as a market-implied forecast of volatility. In practice, forecasters often combine implied volatility with historical estimates to improve predictions (a technique known as volatility forecast combination), but the implied component typically carries a heavy weight.

In this thesis, the forward-looking nature of V2X is especially pertinent. We will be incorporating the V2X as a candidate predictor in volatility forecasting exercises, to benchmark how well a market-based forecast competes with statistical models. The expectation, based on prior evidence, is

---

[62] Poon, S.-H., & Granger, C. W. J. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature, 41*(2), 478–539.

that V2X will prove hard to beat as a 1-month volatility predictor once its bias is accounted for. Indeed, as Carr & Wu (2006) demonstrate, if one uses VIX (or V2X) and adjusts for the average variance risk premium, one gets an almost unbiased forecast of subsequent variance. The informational advantage of implied volatility underscores its importance as a forward-looking risk indicator.

2.3.6 Applications of Implied Volatility: Market Timing and Corporate Finance

Beyond its role in measuring and forecasting volatility, implied volatility (V2X) has practical applications in investment and corporate finance decisions. Because it encapsulates market sentiment and risk expectations, implied volatility can be used for market timing, that is, to inform strategic allocation or entry/exit decisions in asset markets. Investors often view extreme levels of volatility indices as signals. A very high VIX/V2X ("extreme fear") can paradoxically be a contrarian buy signal for equities, on the premise that panic may be near its peak and forward returns are attractive (consistent with the variance risk premium predicting higher returns when it is elevated). Conversely, an ultra-low VIX/V2X ("extreme complacency") might precede market corrections, as it could signal that markets are overlooking risks. While timing the market based on volatility alone is challenging, academic studies have found some merit in volatility-based strategies. As mentioned, a high variance risk premium is associated with higher expected equity returns, so an investor might increase equity exposure when V2X is unusually far above recent realized volatility (indicating fear is high and potential mispricing of risk). Similarly, some volatility traders employ strategies that sell volatility after spikes, banking on mean reversion (these are essentially short-volatility positions that earn the risk premium when volatility eventually falls). However, caution is warranted: volatility spikes are often "short, intense, and not foreseeable", and shorting volatility carries the risk of large losses if another shock hits. Still, implied volatility provides a real-time gauge of market risk appetite, which can supplement other indicators (valuation ratios, credit spreads, etc.) in assessing whether markets are overly optimistic or pessimistic.

In the realm of corporate finance, implied volatility indices are valuable for decision-makers because they reflect the market's uncertainty at any given time. One prominent example is the timing of Initial Public Offerings. Companies generally prefer to go public during periods of market calm and confidence, when valuations are higher and investor demand is robust. High market volatility increases uncertainty about pricing and the risk of a poor reception for a new listing. Empirical observation confirms that when volatility is elevated, IPO activity tends to dry up. Research indicates that during periods when the VIX (and by extension, V2X) is above a threshold (commonly cited is 20 for the VIX, a level associated with heightened volatility), the number of IPOs decreases significantly. This is intuitive: a high "fear index" means investors are risk-averse and markets are turbulent, not an ideal environment for companies seeking to sell new equity. Conversely, when implied volatility remains low for an extended period, signalling stable and favorable conditions, IPO volumes increase markedly. The IPO "window" is considered "open" when volatility indices are low, as companies can fetch better valuations and face less risk of price swings around their debut. The logic extends to other corporate finance activities too: mergers and acquisitions, debt issuance, and share buybacks are all influenced by the level of market volatility.

CFOs monitor indices like VIX/V2X as a barometer of market conditions; a spike in V2X might cause a firm to delay an equity offering or rethink the timing of a major transaction. High implied volatility implies a higher required return from investors (due to risk), which could translate into a lower price if a company issues stock, or a higher interest rate if issuing debt. Thus, companies often wait for volatility to settle before executing financing plans. In our context, Chapter 3 will specifically explore IPO timing and pricing, where we include market volatility (proxied by V2X) as a key explanatory variable. The expectation is that higher V2X will be associated with weaker IPO activity or greater IPO underpricing, consistent with the idea that issuers must compensate investors more in volatile times or avoid going public until the "fear gauge" recedes.

Finally, implied volatility will feature in our volatility forecast evaluation. There, we will use V2X as a benchmark forecast to compare against other models. Its inclusion is motivated by all the reasons discussed: it is a model-free, market-based predictor incorporating information that pure statistical models might miss. By examining how well V2X predicts subsequent realized volatility (relative to, say, a GARCH or historical average), we can evaluate the practical value of implied volatility in forecasting contexts. Moreover, the analysis will shed light on whether the volatility risk premium is time-varying in a way that could be exploitable: for instance, if a simple bias-corrected V2X forecast outperforms other methods, it underscores the efficiency of information in option markets.

Implied volatility is a cornerstone of modern financial analysis for volatility. We have defined it and explained its construction as a model-free, forward-looking measure of expected market volatility. We have seen that it carries rich information: it spikes during crises (earning the "fear index" label), correlates negatively with equity returns contemporaneously, and bears a systematic risk premium reflecting investors' aversion to volatility risk. Implied volatility tends to lead and predict realized volatility to a significant extent, making it a valuable tool for forecasting and risk management. Its applications extend to practical domains like market timing (signalling risk-on/risk-off regimes) and corporate finance decisions (gauging the feasibility of market transactions).

2.4 Variance and Volatility Risk Premia

Implied volatility indices such as the V2X (VSTOXX) are often described as "forecasts" of future market volatility. However, option prices, and hence implied volatility are formed under the risk-neutral probability measure, not under the real-world (physical) measure. As a result, the V2X contains not only expectations about future realized volatility but also a risk premium component reflecting compensation demanded by investors for bearing volatility (variance) risk. This wedge between risk-neutral and physical expectations is central for interpreting implied volatility and for comparing implied measures with purely statistical forecasts of realized volatility.

This section formalizes the concepts of the variance risk premium (VarRP) and the closely related volatility risk premium (VolRP), and it describes how these quantities are constructed empirically using the SX5E and V2X series.

2.4.1 Risk-neutral versus physical expected variance

Let r_t denote the daily log return on the EURO STOXX 50 index (SX5E). Define the (annualized) realized variance over a horizon of H trading days as:

$$RV_{t\to t+H}^{\mathrm{ann}} = \frac{252}{H}\sum_{j=1}^{H} r_{t+j}^2 \,.$$

Option prices, by contrast, embed the risk-neutral expectation of future variance. A model-free implied volatility index like V2X is closely related to the square root of the risk-neutral expected integrated variance over roughly the next month. Let IV_t denote the V2X level expressed in annualized volatility terms (percent). Converting to a decimal volatility gives $iv_t = IV_t/100$. The corresponding implied variance proxy is iv_t^2.

The variance risk premium is commonly defined (in variance units) as the difference between the physical expected variance and the risk-neutral expected variance:

$$VarRP_t \equiv \mathbb{E}_t^{\mathbb{P}}[RV_{t\to t+H}^{\mathrm{ann}}] \; - \; \mathbb{E}_t^{\mathbb{Q}}[RV_{t\to t+H}^{\mathrm{ann}}]\,.$$

In equity index markets, it is typical to find $\mathbb{E}^{\mathbb{Q}}[\cdot]$ above $\mathbb{E}^{\mathbb{P}}[\cdot]$ (investors pay up for crash insurance), which implies $VarRP_t$ is often negative under this definition.

Because $\mathbb{E}_t^{\mathbb{P}}[RV]$ is not directly observable, empirical work typically proxies the premium using realized outcomes (ex post) or trailing realized measures (ex ante), described next.

2.4.2 Empirical construction of implied and realized variance measures

Implied variance proxy. In this thesis we use the V2X level as an observable proxy for risk-neutral expected volatility. In practice we form

$$IVAR_t^{ann} = \left(\frac{V2X_t}{100}\right)^2 ,$$

which is the annualized implied variance in decimal units.

Forward realized variance (ex post). To match the one-month horizon of V2X, we set $H = 30$ trading days and compute:

$$RV_{t\to t+30}^{\mathrm{ann}} = \frac{252}{30}\sum_{j=1}^{30} r_{t+j}^2 \,.$$

This uses future returns from $t + 1$ through $t + 30$ and is therefore an ex post realized benchmark.

Trailing realized variance (ex ante proxy). A trailing realized variance measure can be constructed analogously:

$$RV_{t-30\to t}^{\text{ann}} = \frac{252}{30}\sum_{j=0}^{29} r_{t-j}^2 \, .$$

While backward-looking, this is sometimes used as a proxy for the market's physical expectation of near-term variance when building real-time indicators.

2.4.3 Variance risk premium versus volatility risk premium

There are two closely related "risk premium" objects in the literature:

1. Variance risk premium (variance units).
   A practical ex post proxy consistent with the definition above is:

$$\widehat{VarRP}_t^{(ex\ post)} = RV_{(t\to t+30)}^{(ann)} - IVAR_t^{(ann)} \, .$$

Under this sign convention, equity index data typically yield negative values on average (because implied variance exceeds realized variance).

Some authors reverse the sign and define the premium as implied minus realized:

$$IVAR_t^{\text{ann}} - RV_{t\to t+30}^{\text{ann}} \, ,$$

which is then typically positive. Either convention is fine as long as it is stated clearly and used consistently.

2. Volatility risk premium (volatility points).
   Because V2X itself is quoted in volatility (not variance), practitioners often work with the difference in volatility levels:

$$\widehat{VolRP}_t^{\text{ex post}} = V2X_t - RVOL_{t\to t+30}^{\text{ann}} \, ,$$

$$RVOL_{t\to t+30}^{\text{ann}} = 100\sqrt{RV_{t\to t+30}^{\text{ann}}} \, .$$

This measure is expressed in annualized percentage points ("vol points"). It is intuitive and easy to interpret, but it is important to note that it is not algebraically identical to the variance premium because the square-root mapping is nonlinear.

In this work, the volatility-point premium is particularly useful for interpretation and for connecting to IPO decision-making, since market participants tend to discuss volatility levels (e.g., "V2X is 30") rather than variance.

2.4.4 Economic interpretation and relevance

A persistently positive wedge $V2X_t > RVOL^{ann}_{t\to t+30}$ indicates that options markets price volatility above what subsequently realizes on average. Economically, this reflects the fact that being "short volatility" (selling options or variance swaps) earns an insurance premium most of the time, while occasionally suffering large losses in crisis episodes. Therefore, the premium can be interpreted as compensation for bearing tail risk and for providing crash insurance.

For the purposes of this thesis, this matters for two reasons:

- Interpreting implied volatility as a forecast. Because V2X is risk-neutral and includes a premium, it is not a pure forecast of realized volatility under the physical measure. When we compare V2X-based forecasts with statistical forecasts (rolling, EWMA, GARCH-family), we must keep in mind that differences may reflect risk premia, not only forecasting "skill."
- Measuring the financing environment beyond volatility. The premium component is closely related to risk aversion and market stress, which may influence corporate financing decisions independently of expected realized volatility. In Chapter 3, volatility forecasts proxy market uncertainty faced by issuers; the risk premium component can be interpreted as capturing the price investors place on uncertainty at that moment (i.e., how costly it is to raise capital when protection is expensive).

Empirically, we construct both a forward-looking realized volatility measure (30 trading days ahead) and the corresponding volatility risk premium:

$$\widehat{VolRP}_t^{\text{ex post}} = V2X_t - RVOL^{\text{ann}}_{t\to t+30}\,.$$

2.5 GARCH Family Models

ARCH/GARCH models provide a parsimonious and widely used framework for modelling the conditional variance of financial returns. Unlike rolling or EWMA estimators, which are purely mechanical summaries of past returns, GARCH-family models specify a parametric data-generating process for volatility that can be estimated via maximum likelihood and used to produce model-based volatility forecasts at different horizons. These models are especially suited to the stylized facts documented earlier, most importantly volatility clustering and persistence, and they form a standard benchmark class in volatility forecasting.

2.5.1 General Setup

Let P_t be the EURO STOXX 50 index level (SX5E) and define the daily log return:

$$r_t = \ln(P_t) - \ln(P_{t-1}) \ .$$

We specify a conditional mean–variance decomposition:

$r_t = \mu + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t,$

Where $\sigma_t^2 = \mathrm{Var}(\, \varepsilon_t \mid \mathcal{F}_{t-1}\,)$ is the conditional variance and z_t is an i.i.d. innovation with mean zero and unit variance. In practice, financial returns often exhibit excess kurtosis ("fat tails"), therefore we allow $z_t$to follow either a Gaussian distribution or a heavy-tailed Student-t distribution. The latter is often empirically preferred for equity index returns.

The core modelling task is to specify a recursion for σ_t^2as a function of past shocks and past variances, and then to estimate the parameters by maximum likelihood.

2.5.2 The GARCH(1,1) model

The benchmark specification is the *GARCH(1,1)* model:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \,,$$

$$\omega > 0, \alpha \geq 0, \beta \geq 0 \,.$$

Here, α captures the short-run impact of new shocks on volatility, while βcaptures persistence through the lagged conditional variance. The sum $\alpha + \beta$ is a convenient measure of *volatility persistence*: values close to 1 imply that volatility shocks decay slowly, consistent with volatility clustering. Under $\alpha + \beta < 1$, the process is covariance-stationary with unconditional variance $\omega/(1 - \alpha - \beta)$.

2.5.3 Asymmetric GARCH models: leverage effects

A limitation of symmetric GARCH is that it treats positive and negative return shocks of equal magnitude as having the same effect on future volatility. For equity markets, this is often unrealistic because negative returns tend to increase future volatility more strongly (the "leverage effect"). Two widely used asymmetric extensions are:

(i) GJR-GARCH (Threshold GARCH).

The *GJR-GARCH(1,1)* model adds an indicator for negative shocks:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{(t-1)}^2 + \gamma 1_{(\varepsilon_{(t-1)}<0)}\varepsilon_{(t-1)}^2 + \beta\sigma_{(t-1)}^2 \, .$$

(ii) EGARCH (Exponential GARCH).

The EGARCH(1,1) model, introduced by Nelson (1991), instead models the logarithm of the conditional variance:

$$\log(\sigma_t^2) = \omega + \beta \log(\sigma_{t-1}^2) + \alpha \left|\frac{\varepsilon_{t-1}}{\sigma_{t-1}}\right| + \gamma \frac{\varepsilon_{t-1}}{\sigma_{t-1}}$$

Because variance is modelled in logs, EGARCH does not require non-negativity constraints on parameters, and the leverage term γ captures asymmetry: $\gamma < 0$ typically corresponds to negative shocks increasing volatility more strongly than positive shocks.

Model	α	β	Asymmetry term γ	ν(t df)	LogLik	AIC	BIC	Persistence
GARCH(1,1)	0.1108	0.8782	—	5.7550	−7708.51	15427.0	15459.7	$\alpha + \beta = 0.9890$
GJR-GARCH(1,1)	~0.0000	0.8803	0.2052	6.5470	−7600.13	15212.3	15251.5	$\alpha + \beta + \gamma/2 = 0.9829$
EGARCH(1,1)	0.1186	0.9766	−0.1718	6.7392	−7570.01	15152.0	15191.3	$\beta \approx 0.9766$

Table 2.2 reports maximum likelihood estimates for the GARCH(1,1), GJR-GARCH(1,1), and EGARCH(1,1) specifications using Student-t innovations. Across models, volatility is highly persistent (persistence measures close to one), consistent with volatility clustering in equity index returns. Both asymmetric specifications indicate a statistically meaningful leverage effect: negative shocks are associated with larger increases in conditional volatility. Information criteria favour EGARCH in-sample (lowest AIC/BIC), followed by GJR-GARCH and then symmetric GARCH.

2.5.4 Forecasting with GARCH-family models

Once the parameters are estimated, the models yield conditional variance forecasts. The one-step-ahead forecast is $\hat{\sigma}_{t+1|t}^2$. Multi-step forecasts $\hat{\sigma}_{t+h|t}^2$ for $h = 1, \ldots, H$ can be obtained recursively using the model dynamics.

To align with the 30-day horizon commonly associated with implied volatility indices, we also construct a *30-trading-day volatility forecast* from the path of forecasted daily variances:

$$\hat{\sigma}_{t,H}^{ann} = 100\sqrt{\left(\frac{252}{H}\right)\sum_{h=1}^{H}\hat{\sigma}^2(t+h|t)}\ ,$$

where $H = 30$ trading days and the factor $252/H$ annualizes the average predicted daily variance over the horizon. This produces forecasts in *annualized volatility percentage points*, directly comparable across model classes and to option-implied measures.

Using the estimated models, we generate out-of-sample volatility forecasts over 2010-2024. For each trading day t, we compute both a one-day-ahead annualized volatility forecast and a 30-trading-day annualized forecast constructed from the model-implied variance path. These forecast series are saved and later evaluated against realized volatility benchmarks and matched to IPO event dates.

In the empirical analysis, we estimate and compare the forecasting performance of GARCH(1,1), GJR-GARCH(1,1), and EGARCH(1,1). The models are fit by maximum likelihood using daily SX5E returns. Forecasts are generated out-of-sample and evaluated using standard loss functions and statistical tests. In addition, because forecast horizons and the timing of information matter for IPO decisions, forecasts are computed using only information available at time t and then matched to IPO event dates.

2.6 Realized Volatility Construction

Although volatility is latent, the empirical analysis requires an *observable* benchmark against which volatility forecasts can be evaluated. This thesis therefore constructs realized volatility measures from the SX5E daily return series and uses them as (i) an ex-post benchmark for forecast evaluation, and (ii) an input for the variance/volatility risk premium measures discussed in Section 2.4.

2.6.1 Realized variance and annualised realized volatility

Let r_t denote daily log returns. For a horizon of H trading days, the realized variance over a window is computed as the average of squared daily returns, annualised using 252 trading days:

$$RV_{t,H}^{\text{ann}} = \frac{252}{H}\sum_{j=0}^{H-1} r_{t-j}^2 \ .$$

In the baseline analysis, we set $H = 30$ trading days to remain consistent with the "monthly" horizon used throughout the chapter (and to keep the realized-volatility benchmark aligned with the 30-day focus of implied-volatility indices such as V2X).

2.6.2 Trailing versus forward realized volatility

It is important to distinguish two realized-volatility objects that play different roles in the thesis:

(i) Trailing realized volatility (ex-ante proxy).
The trailing measure uses returns observed up to time t (backward-looking):

$$RVOL_{t,30}^{\mathrm{ann}}(\mathrm{trailing}) = 100\sqrt{\frac{252}{30}\sum_{j=0}^{29} r_{t-j}^2}\,.$$

(ii) Forward realized volatility (ex-post benchmark).
The forward measure uses returns observed after time t (forward-looking):

$$\mathrm{RVOL}_{t,30}^{\mathrm{ann}}(\mathrm{forward}) = 100\sqrt{\frac{252}{30}\sum_{j=1}^{30} r_{t+j}^2}$$

This forward measure is the natural target when we compare model forecasts formed at time t with the subsequent volatility outcome in Section 2.7. Because it requires future returns, it is observable only ex post and is never used as an input to any forecast, it serves exclusively as the evaluation benchmark.

Note on overlap with rolling volatility.
When trailing realized volatility is computed from squared returns over the last 30 trading days, it is (up to small numerical differences and the return-mean assumption) effectively the same object as a 30-day rolling close-to-close volatility estimate. For this reason, the thesis keeps the realized-volatility notation mainly to make the distinction between *trailing* and *forward* windows explicit, and to maintain consistency with the risk premium definitions in Section 2.4.

2.6.3 Practical alignment and sample availability

The realized-volatility measures are computed on the SX5E trading-day calendar. Using $H = 30$ implies:

- the trailing series starts after the first 30 return observations;
- the forward series is unavailable for the final 30 trading days of the sample (since future returns are required).

After construction, the realized-volatility series are merged with V2X and with the competing volatility forecasts (historical estimators and GARCH-family forecasts). The resulting aligned dataset is then used (i) to evaluate forecast performance using loss functions and statistical tests (Section 2.7), and (ii) to match market uncertainty measures to IPO event dates in Chapter 3.

2.7 Forecast Evaluation Framework

Before linking volatility to IPO outcomes in Chapter 3, it is necessary to verify that the volatility measures used in the empirical models behave as credible real-time proxies for market uncertainty. Since volatility is latent, forecast evaluation requires (i) an ex post benchmark, typically realized volatility constructed from returns, and (ii) an out-of-sample (OOS) design that mimics the information set available at the time forecasts are formed.

In this section we evaluate the predictive performance of the main volatility measures considered in this thesis—historical estimators (Rolling(30), EWMA(0.94)), implied volatility (V2X), and GARCH-family forecasts—against a forward realized volatility benchmark at a one-month horizon.

2.7.1 Forecast targets and horizon alignment

Let $H = 30$ denote the forecast horizon in trading days. The benchmark target is the 30-day forward realized volatility, annualised:

$$RV_{t\to t+H}^{\text{ann}} = 100\sqrt{\frac{252}{H}\sum_{j=1}^{H} r_{t+j}^{2}}\,, H = 30\,.$$

All competing forecasts are expressed in *annualized volatility percentage points*, so that forecast errors are directly interpretable in "vol points".

Forecast evaluation is conducted *out-of-sample* over 2010–2024, consistent with the IPO dataset start date in Chapter 3. The resulting OOS evaluation sample contains $N = 3793$ daily observations (excluding the final Hdays required to compute forward realized volatility).

2.7.2 Correlation structure across volatility measures

As a preliminary diagnostic, it is useful to assess how strongly the different volatility measures co-move. Figure 2.9 reports the correlation matrix for implied volatility (V2X), trailing and forward realized volatility, historical estimators, and 30-day GARCH-family forecasts over the OOS period.

The correlation matrix shows three main patterns. First, the *historical measures are extremely close to one another*: trailing realized volatility and Rolling(30) are mechanically equivalent in construction and therefore exhibit a correlation of 1.00, while EWMA(0.94) remains very highly correlated with both (about 0.97). Second, implied volatility is strongly related to current/historical volatility regimes: V2X correlates about 0.82 with trailing realized volatility, 0.81 with Rolling(30), and 0.87 with EWMA. Third, correlations with *forward realized volatility are noticeably lower*, reflecting the intrinsic uncertainty in predicting future volatility: forward RV correlates about 0.60 with V2X and about 0.52–0.55 with the historical measures. Finally, the GARCH-family 30-day forecasts are highly correlated with one another (roughly 0.90–0.95) and correlate strongly with trailing measures (around 0.83–0.90), but less with forward RV (0.40–0.44), indicating that these models mostly track persistent volatility regimes rather than capturing all variation in future realized volatility.

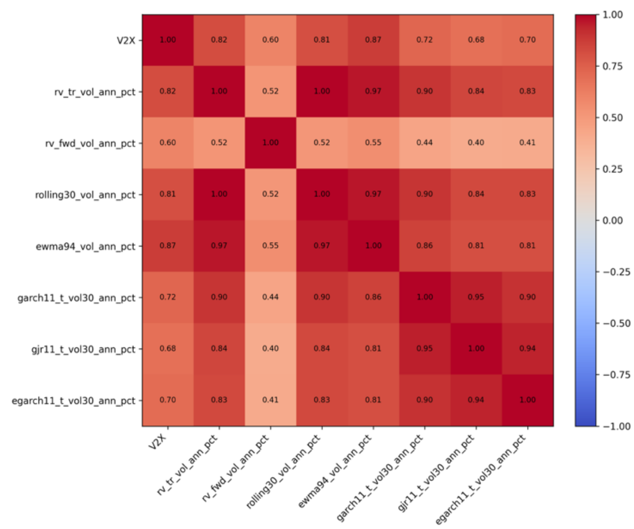

Figure 2.9: Correlation matrix of implied volatility (V2X), realized volatility measures, historical estimators, and 30-day GARCH-family forecasts (OOS sample 2010–2024).

2.7.3 Loss functions

Forecast accuracy is evaluated using three complementary loss functions. Let $\hat{\sigma}_t^{(m)}(H)$ denote model m's annualised volatility forecast for horizon H, and let $RV_{t\to t+H}^{\text{ann}}$ denote the realized outcome.

(i) Root Mean Squared Error (RMSE)

$$RMSE_m = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\widehat{\sigma_t^{(m)}}(H) - RV_{t\to t+H}^{\text{ann}}\right)^2}\ .$$

RMSE penalizes large errors more strongly and is widely reported in volatility forecast studies.

(ii) Mean Absolute Error (MAE)

$$MAE_m = \frac{1}{T}\sum_{t=1}^{T}\left|\widehat{\sigma_t^{(m)}}(H) - RV_{t\to t+H}^{\mathrm{ann}}\right|.$$

MAE provides a more robust measure of typical forecast error magnitude in volatility units (percentage points).

(iii) QLIKE (Quasi-likelihood loss)
Because most econometric models naturally generate variance forecasts and because volatility is latent, a variance-based loss is also reported. Define annualized variance forecasts as $\hat{v}_t^{(m)} = (\hat{\sigma}_t^{(m)}/100)^2$and realised annualised variance as $RV_{t\to t+H}^2 = (RV_{t\to t+H}^{\mathrm{ann}}/100)^2$. QLIKE is:

$$QLIKE_m = \frac{1}{T}\sum_{t=1}^{T}\left[\log\left(\widehat{v_t^{(m)}}\right) + \frac{RV_{t\to t+H}^2}{\widehat{v_t^{(m)}}}\right].$$

QLIKE is particularly common in volatility forecasting because it is closely related to the Gaussian quasi-likelihood and tends to be less sensitive to extreme volatility realizations than squared-error loss.

Model	N	RMSE (vol, pp)	MAE (vol, pp)	QLIKE (var)
V2X	3793	7.978975	5.957836	-2.299102
EWMA(0.94)	3793	7.992172	5.307039	-2.267306
Rolling30	3793	8.440045	5.566164	-2.229602
GARCH(1,1)-t	3793	8.746562	5.872794	-2.224400
EGARCH(1,1)-t	3793	8.503092	6.000288	-2.202755
GJR-GARCH(1,1)-t	3793	9.491626	6.152888	-2.182339

Table 2.3: Out-of-sample forecast accuracy (H = 30 trading days, 2010–2024).

This summarizes OOS forecast accuracy for the 30-day horizon using RMSE, MAE, and QLIKE. Overall, V2X performs best according to RMSE and QLIKE: its RMSE is 7.98 vol points and its QLIKE is −2.299. EWMA(0.94) is extremely close in RMSE (7.99) and achieves the lowest MAE (5.31), implying slightly smaller typical absolute errors even though V2X performs marginally better under RMSE/QLIKE. Rolling(30) is weaker than both V2X and EWMA (RMSE 8.44, QLIKE −2.230), consistent with the idea that equal-weight rolling windows can lag at turning points. The GARCH-family forecasts do not outperform implied volatility in this sample: GARCH(1,1)-t (RMSE 8.75) and EGARCH(1,1)-t (RMSE 8.50) are close but still behind, while GJR-GARCH(1,1)-t performs worst (RMSE 9.49, QLIKE −2.182). Taken together, these results suggest that (i) implied volatility contains substantial forward-looking information for the one-month horizon, and (ii) simple historical measures remain competitive benchmarks, especially EWMA.

2.7.4 Statistical comparison: Diebold–Mariano tests

Differences in average loss values may not be statistically meaningful. To test whether forecast performance differs significantly, we apply Diebold–Mariano (DM) tests of equal predictive accuracy. For two forecasts $a$and b, define the loss differential $d_t = L_t^{(a)} - L_t^{(b)}$. The null hypothesis is $H_0: \mathbb{E}[d_t] = 0$.

Because the target is a 30-day forward realized volatility, the realized benchmark is constructed from overlapping windows, generating serial correlation in loss differentials. We therefore compute DM tests using HAC (Newey–West) standard errors with $H - 1 = 29$ lags.

Table 2.4 reports DM tests using QLIKE losses, benchmarked against V2X. Mean differentials are positive for all competing models, indicating higher average loss than V2X. The difference is statistically significant for Rolling(30) (DM $t = 2.30$, $p = 0.021$), implying that V2X outperforms Rolling(30) in a statistically meaningful way under QLIKE. Differences between V2X and the GARCH-family forecasts are weaker: EGARCH(1,1)-t is borderline at the 5% level ($p = 0.051$) and GJR-GARCH(1,1)-t is borderline at the 10% level ($p = 0.062$), while standard GARCH(1,1)-t ($p = 0.151$) and EWMA(0.94) ($p = 0.194$) are not significantly different from V2X. This pattern reinforces the main conclusion from the table: implied volatility is difficult to beat at the one-month horizon, while EWMA provides a very competitive baseline.

Model A	Model B	N	Mean diff (A—B)	DM t-stat	p-value
Rolling30	V2X	3793	0.069500	2.302233	0.021322
EGARCH(1,1)-t	V2X	3793	0.096347	1.947321	0.051496
GJR-GARCH(1,1)-t	V2X	3793	0.116763	1.867062	0.061893
GARCH(1,1)-t	V2X	3793	0.074701	1.437322	0.150627
EWMA(0.94)	V2X	3793	0.031796	1.298154	0.194234

Table 2.4: Diebold–Mariano tests (QLIKE loss), benchmark = V2X; HAC standard errors with 29 lags

2.7.5 Implications

The evaluation results motivate two implementation choices for the IPO analysis. First, implied volatility (V2X) is a natural candidate for the baseline "financing environment uncertainty" proxy, given that it provides the best overall forecast performance under RMSE and QLIKE and statistically outperforms Rolling(30). Second, because EWMA is extremely competitive (and not statistically distinguishable from V2X under DM testing), and because GARCH-family forecasts capture persistence and asymmetry, Chapter 3 will treat EWMA and GARCH-family forecasts as important robustness alternatives. This structure ensures that IPO outcome results are not driven by a single volatility proxy and that both market-implied and model-based views of uncertainty are represented.

3 Empirical Analysis Part II: Volatility and IPO Outcomes

Chapter 2 constructed a set of daily volatility measures for the European equity market—ranging from simple historical estimators (rolling and EWMA) to forward-looking implied volatility (V2X) and conditional variance forecasts from GARCH-family models. The purpose of Chapter 3 is to translate those volatility measures into *corporate finance implications* by asking a concrete question: when market uncertainty is higher (as proxied by forecasted volatility), do IPO attempts become more likely to fail or to proceed under worse terms?

This chapter therefore shifts from "which volatility model forecasts best?" (a statistical forecasting question) to "which volatility signal matters for financing outcomes?" (a corporate finance question). The empirical design follows the logic of the IPO process as a sequence of decisions and constraints. Before a stock starts trading, issuers and underwriters face (i) *offer price risk* (uncertainty about the feasible offer price), (ii) *distribution risk* (uncertainty about demand), and (iii) *aftermarket risk* (uncertainty about initial trading performance). A high-volatility environment amplifies all three: it widens the set of plausible valuations, complicates book-building, increases underwriting inventory risk, and raises the probability of a weak aftermarket. In practice, these pressures can lead either to (a) abandoning the transaction (withdrawal), (b) delaying it (postponement), or (c) completing it but at a lower offer price than initially targeted and/or with higher underpricing.

Consistent with this mechanism, Chapter 3 is organized around three IPO outcome dimensions, each matched to a specific step in the issuance process:

1. *Disruption outcome* (failure to complete on schedule): We treat withdrawn and postponed IPOs as *disruption events* (a single binary outcome). The economic interpretation is that the issuer did not successfully transition from the filing/marketing phase to a completed public listing in the intended timeframe. This outcome is the closest empirical proxy—available in the data—for "IPO market timing under uncertainty."
2. *Offer price setting* (revision relative to the expected range): For IPOs that proceed to pricing, we measure whether the final offer price ends up below, within, or above the *expected price range* disclosed earlier in the process (from Bloomberg's filing terms). This captures how market uncertainty translates into offer-price concessions at the pricing stage.
3. *Immediate aftermarket performance* (initial return / underpricing): Because a full event study requires daily post-IPO prices for every firm (which is not consistently available in our current dataset), we focus on the first trading day outcomes that Bloomberg provides directly: Offer-to-1st Open and Offer-to-1st Close returns. These are the cleanest, widely-used underpricing measures and align well with the "aftermarket risk" component.

Conceptually, the volatility measures from Chapter 2 enter Chapter 3 as a *real-time financing environment uncertainty signal*. The econometric analysis uses that signal to explain cross-sectional and time-series variation in disruption probability, pricing outcomes, and initial returns, controlling for issuer characteristics and market structure (country, sector, year effects, and deal size).

A key modelling choice in this chapter is the definition of the relevant *event date* for volatility matching. For successfully completed IPOs ("Trading"), the economically relevant decision point is the pricing (or, when pricing date is unavailable, the listing/effective date). For disrupted IPO attempts, the relevant date is the withdrawal or postponement date—the moment at which the issuer/underwriter effectively decides to stop (withdraw) or delay (postpone) the offering, presumably in response to market conditions.

3.1 IPO Dataset Construction

This section describes the construction of the IPO dataset used in the second part of the empirical analysis as before estimating any models, it is essential to define the underlying unit of observation, the scope of the sample, the outcome definitions, and the data-cleaning choices that produce a consistent and replicable IPO dataset.

3.1.1 Data source, scope, and unit of observation

The IPO data are extracted from Bloomberg's IPO Deal List, covering the period 2010–2024 with a Western European scope. The key feature of using Bloomberg's deal list (rather than only exchange-listed firms) is that the dataset captures IPO attempts, including offerings that do not complete. This is critical for the thesis because the core mechanism is not only that volatility affects pricing, but also that volatility can prevent issuance from occurring at all. In other words, volatility matters for the extensive margin of financing (whether the deal proceeds) as well as the intensive margin (how the deal is priced when it proceeds).

Accordingly, the unit of observation in this chapter is an IPO attempt (deal). Each deal contains identifying information (issuer name, ticker, country/region, primary exchange), classification variables (industry sector, currency), and deal-process variables (status/stage, event dates, offer size, offer price range, offer price, and first-day return measures where available).

A key choice in this thesis is to treat both Withdrawn and Postponed deals as financing disruption events, rather than focusing only on withdrawals. This is consistent with the motivation of Chapter 1: when market uncertainty rises, firms may either abandon an offering or delay it, but in both cases the firm fails to raise equity capital at the intended time and terms. Moreover, postponement is a common response in European IPO markets, so restricting attention to withdrawals alone risks understating the effect of volatility on issuance feasibility.

The "Offer Stage" variable is used to define three mutually exclusive outcomes in the raw extract:

- Trading: IPO successfully priced and begins trading (completed deal)
- Withdrawn: IPO attempt cancelled/withdrawn
- Postponed: IPO attempt delayed/suspended rather than completed as planned

These categories exhaust the states observed in the extracted dataset. On this basis, define:

- $Disruption_i = \mathbb{1}\{OfferStage \in \{Withdrawn, Postponed\}\}$

This disruption indicator becomes the dependent variable for the baseline "disruption risk" model, while pricing and first-day performance outcomes are studied conditional on Trading deals.

3.1.2 Sample size and composition

The raw Bloomberg Terminal extract contains 3,341 IPO attempts over 2010–2024. Sample composition by outcome stage is as follows:

- Trading: 2,953 deals (88.39%)
- Withdrawn: 318 deals (9.52%)
- Postponed: 70 deals (2.10%)

In total, the disruption category contains 388 deals, implying an overall disruption rate of 11.61% in the raw sample (Withdrawn + Postponed). These counts confirm that restricting the analysis to completed IPOs would remove a non-trivial part of the IPO process, exactly the part that is expected to be most sensitive to volatility and market uncertainty.

Coverage of the date variable is substantial but not complete:

- Disruption events: 388 total
- Usable disruption date available: 275 deals
- Missing disruption date: 113 deals
- Date coverage among disruptions: 70.88% (275/388)

This distinction is central for the empirical design. The raw sample can be used for descriptive statistics, but models that require volatility at the disruption decision date must use the subsample of disruptions with valid dates.

3.1.3 Time distribution of IPO attempts (2010–2024)

IPO activity varies substantially over time.

ann_year	postponed	trading	withdrawn
2010	0	154	37
2011	0	133	39
2012	0	90	22
2013	0	146	20
2014	0	256	48
2015	9	263	38
2016	4	214	26

2017	6	294	10
2018	21	239	23
2019	8	139	9
2020	2	192	6
2021	12	548	24
2022	6	138	7
2023	0	78	4
2024	2	69	5

Table 3.1 reports counts by announcement year and Offer Stage

To describe the pipeline of issuance attempts, this thesis summarizes IPO counts by announcement year. The year-by-stage breakdown shows large differences in both overall activity and disruption frequency.

A striking feature of the sample is the concentration of activity in certain years. For example, 2021 is the largest year in the dataset with 584 total IPO attempts, of which 548 reach Trading and 36 are disrupted (withdrawn or postponed). Activity is also relatively high in the mid-2010s (e.g., 2014–2018), while the final years of the sample show substantially fewer attempts (e.g., 76 total attempts in 2024).

Disruption rates also vary over time. Early years (2010–2012) show disruption rates around ∼19–23% in this dataset, whereas several later years show much lower disruption rates (for example, 2017 is notably low in the announcement-year summary). These patterns motivate the inclusion of year fixed effects in the regression analysis and highlight why it is necessary to match each IPO to the correct contemporaneous volatility environment rather than relying on pooled averages.

3.1.4 Cross-sectional distribution: countries and sectors

The Bloomberg extract spans 33 country/region codes and 12 industry-sector categories. The distribution is concentrated: the five largest country codes account for approximately 69.7% of all IPO attempts in the sample. In particular, the UK-coded group (GB) is the single largest country category in the dataset (944 total attempts), followed by several other major Western European IPO markets. The remaining country groups are smaller but still important for identifying cross-country variation in IPO outcomes and institutional settings.

Country/Region	postponed	trading	withdrawn
GB	16	819	109
SW	3	525	20
IT	5	352	25
FR	8	224	20
GE	10	169	25

NO	4	163	9
FI	0	97	10
DE	2	94	5
GS	3	75	3
SZ	2	71	12
NE	2	65	19
SP	4	62	9
IR	2	44	6
LX	2	36	11
BE	2	32	2
JE	0	28	7
GR	1	19	11
other countries	4	76	14

Table 3.2 reports counts by Country/Region code and Offer Stage

On the sector dimension, IPO attempts are also concentrated. The five largest industry sectors account for approximately 78.9% of all IPO attempts. The largest sectors include Consumer, Non-cyclical, Financial, Technology, Industrial, and Consumer, Cyclical. Sector-level disruption rates are not uniform: some sectors appear systematically more disruption-prone than others (for example, Financial and Industrial have sizeable disruption counts), although sector differences must ultimately be interpreted cautiously because the sector categories aggregate heterogeneous firms and deal types.

Industry Sector	postponed	trading	withdrawn
Consumer, Non-cyclical	10	684	68
Financial	29	570	62
Technology	5	388	26
Industrial	5	361	47
Consumer, Cyclical	8	332	40
Communications	5	232	23
Energy	4	180	19
Basic Materials	2	104	20
Diversified	0	72	3
Utilities	2	23	5
Government	0	3	0
Funds	0	2	1

Table 3.3 reports counts by Industry Sector and Offer Stage

These descriptive distributions matter for two reasons:

1. Model specification: Country and sector fixed effects help control for persistent heterogeneity in IPO markets (regulatory environments, investor bases, typical deal sizes, valuation norms).
2. Interpretation: If volatility affects sectors differently (e.g., growth sectors vs defensive sectors), this can be examined later through interaction terms or subsample tests.

3.2 Matching IPOs to Volatility Forecasts

Section 3.1 established the IPO attempt universe and defined the key financing outcomes (Trading versus disruptions). This section connects the deal-level IPO dataset to the daily volatility measures and forecasts developed in Chapter 2, creating the analysis-ready panel that underpins the empirical models in Sections 3.3 and 3.4. Because volatility is measured as a trading-day time series while IPO events are recorded as calendar dates, the matching strategy must (i) respect the timing of information availability (avoid look-ahead bias), (ii) define a consistent IPO "decision date" across outcomes, and (iii) ensure that volatility values assigned to IPOs are contemporaneous rather than stale.

Chapter 2 constructed multiple daily measures of market uncertainty for the European equity market. These measures are interpreted here as *financing environment uncertainty*, a state variable summarising how difficult it is for issuers and underwriters to place equity at predictable terms under demand and aftermarket risk.

The volatility dataset used for matching is a daily panel (trading-day frequency) containing, for each date t, the following volatility proxies (expressed as annualised percentages unless noted):

- V2X: option-implied volatility (forward-looking; index points, approximately a 30-day horizon).
- Historical estimators:
    - Rolling 30-day volatility (based on SX5E returns),
    - EWMA volatility with λ=0.94.
- GARCH-family forecasts (30-day horizon):
    - GARCH(1,1)-t,
    - GJR-GARCH(1,1)-t,
    - EGARCH(1,1)-t.

A key advantage of this design is that it does not rely on a single measure of uncertainty. Implied volatility reflects forward expectations embedded in option prices; historical measures provide transparent backward-looking benchmarks; and GARCH forecasts capture conditional volatility dynamics. Using several proxies allows the empirical analysis to test whether the financing implications are robust across different sources of volatility information.

3.2.1 Defining the IPO decision date

The central matching problem is choosing the correct date at which volatility should be measured for each IPO attempt. In corporate finance terms, volatility should be measured when the IPO outcome is resolved, so when the issuer either proceeds to price the deal or decides to disrupt it.

To implement this consistently, the thesis defines an IPO-specific decision date:

- Trading deals: the decision date is the Pricing Date (bookbuilding resolves into the offer price). If Pricing Date is missing, the Effective Date is used as a fallback (typically close to the first trading day and still tightly linked to completion).
- Disruption deals (Withdrawn or Postponed): the decision date is the Withdrawal/Postponement date

This alignment matters because it matches the volatility measure to the economic mechanism emphasized in Chapter 1: firms tend to withdraw or delay when market uncertainty spikes near the point of execution, not necessarily at filing. Measuring volatility at the resolution date is therefore the most direct way to capture the uncertainty the issuer faced when the final decision was made.

3.2.2 Trading-day alignment and no look-ahead matching

The volatility dataset is defined on trading days, while IPO decision dates can fall on weekends/holidays or on days with missing volatility observations. To avoid look-ahead bias, each IPO attempt is matched to volatility using an as-of merge:

$$\text{vol\_date}_i = \max\{\, t \le \text{decision\_date}_i \},$$

i.e., the most recent trading day on or before the IPO decision date. This ensures that the volatility value assigned to an IPO is based only on information that would have been observable at or before the decision date.

A practical diagnostic for match quality is the lag:

$$\text{match\_lag\_days}_i = \text{decision\_date}_i - \text{vol\_date}_i.$$

If the volatility dataset fully covers the IPO period, lags should be near zero (0–2 days) reflecting weekends/holidays. If the volatility dataset ends before the IPO dataset ends, the merge will mechanically attach the last available volatility date to later IPOs, creating stale matches with large lags. Those observations must be excluded to avoid mismeasuring uncertainty.

3.2.3 Matching diagnostics and resulting estimation sample

Starting from the 3,341 IPO attempts in the raw Bloomberg extract (Section 3.1), the decision-date requirement and the volatility match produce the following sample evolution:

1. Decision date available: 3,232 IPO attempts have a usable decision date. The only meaningful loss at this stage comes from disruptions with missing manually collected disruption dates: 109 disruption events lack a usable decision date and therefore cannot be used in event-date-based regressions.
2. Volatility matching: all 3,232 deals can be matched mechanically to the volatility panel via the as-of merge. However, the initial lag diagnostics show a small number of extremely large lags (maximum lag 358 days) a clear sign of stale matches occurring when IPO decision dates extend beyond the available volatility sample.
3. Dropping stale matches: after excluding IPO attempts with a lag larger than 5 days, the final matched dataset contains 3,217 observations. In other words, only 15 deals are removed by the stale-match filter. This indicates that the volatility dataset covers almost the entire IPO decision-date period, with a small loss concentrated at the very end of the sample.

The resulting matched dataset used for Model 1 (disruption risk) spans decision dates from 2010-01-11 to 2024-11-12 (reflecting the overlap between IPO decision dates and the volatility panel).

After applying the decision-date requirement and stale-match filter, the final dataset contains:

- Trading: 2,942 deals
- Withdrawn: 208 deals
- Postponed: 67 deals
- Total disruptions (Withdrawn + Postponed): 275 deals

This implies an event-date-based disruption rate of:

$$\hat{P}(Disruption = 1) = \frac{275}{3217} = 0.0855$$

i.e., 8.55%. This rate is lower than the raw disruption rate in Section 3.1 because disruptions without a usable disruption date are excluded from event-date matching by construction.

After removing stale matches, match lags are extremely small:

- Max lag: 3 days
- Median lag: 0 days
- Mean lag: 0.009 days
- Lag distribution:
    - 0 days: 3,200 deals (99.47%)

- 1 day: 6 deals (0.19%)
- 2 days: 10 deals (0.31%)
- 3 days: 1 deal (0.03%)

This provides strong reassurance that the volatility variables assigned to IPO attempts are contemporaneous (same trading day in almost all cases) and not driven by mechanical matching artifacts.

3.2.4 Standardizing volatility proxies for interpretation

Because the level and dispersion differ across volatility proxies (e.g., V2X index points vs model-based volatility in percent), each volatility proxy Vol is transformed into a z-score:

$$zVol_i = \frac{Vol_i - \overline{Vol}}{s(Vol)}.$$

This standardization makes coefficients comparable across models: in the regressions that follow, the key coefficient is interpretable as the effect of a one-standard-deviation increase in financing-environment uncertainty on the probability of disruption or on pricing outcomes.

3.3 Empirical Model 1: IPO Disruption (Withdrawn or Postponed)

Section 3.1 established the universe of Western European IPO attempts and defined "disruption" events as deals that are either Withdrawn or Postponed. Section 3.2 then matched each IPO attempt to the daily volatility measures developed in Chapter 2 at an IPO-specific decision date (pricing date for completed deals and the manually collected disruption date for disrupted deals), ensuring no look-ahead bias and eliminating stale matches. This section uses the resulting matched dataset to estimate whether higher financing-environment uncertainty, proxied by forecasted volatility, a higher probability that an IPO attempt is disrupted rather than successfully completed.

This model is the empirical counterpart of the work's core corporate-finance mechanism. When volatility rises, the uncertainty around the feasible offer price increases, book-building demand becomes harder to forecast, and underwriters face greater distribution and inventory risk. In such states, issuers may rationally choose to delay the offering (postpone) or abandon it (withdraw), especially if they perceive that the market cannot clear the deal at acceptable terms. Because this decision occurs before (or instead of) pricing, the disruption outcome is the most direct observable manifestation of volatility affecting financing feasibility.

3.3.1 Dependent variable and estimation sample

The dependent variable is a binary indicator:

$$Disruption_i = \mathbb{1}\{\text{Offer Stage}_i \in \{\text{Withdrawn, Postponed}\}\}.$$

In contrast to studies that focus only on successful IPOs (and thus can only speak to pricing outcomes conditional on completion), this specification explicitly models the extensive margin of going public. This is important in the European context where "postponement" is common and represents a meaningful financing failure at the planned time even if the firm may return later.

The estimation sample is the matched dataset from Section 3.2 after removing stale volatility matches (lag $>$ 5 days). It contains:

- N = 3,217 IPO attempts with valid decision dates and contemporaneous volatility matches;
- A disruption rate of 8.55% (Withdrawn + Postponed), lower than the raw disruption rate because some disrupted deals lacked a usable manually collected disruption date and therefore could not be used in event-date matching.

This sample forms the baseline for Model 1.

3.3.2 Key explanatory variables: volatility proxies from Chapter 2

The key regressors are alternative volatility proxies measured at the IPO decision date. Chapter 2 produced several model classes that represent different information sets:

1. V2X implied volatility (forward-looking; option-implied market expectation of near-term volatility);
2. Historical volatility benchmarks (rolling 30-day volatility and EWMA λ=0.94);
3. GARCH-family forecasts (GARCH, GJR-GARCH, EGARCH).

These measures are not inserted simultaneously in one regression because they are highly correlated and would create collinearity and unstable inference. Instead, we estimates separate regressions, one per volatility proxy, so each model answers: *does this specific uncertainty measure predict disruption risk?*

To make coefficients comparable across proxies, each volatility series is standardized:

$$zVol_i = \frac{Vol_i - \overline{Vol}}{s(Vol)}.$$

Thus, the coefficient on $zVol$ corresponds to a one standard deviation increase in the volatility proxy.

3.3.3 Econometric specification

The disruption outcome is modelled using a binomial generalized linear model (logit link):

$$P(Disruption_i = 1) = \Lambda\big(\alpha + \beta zVol_i + \gamma' X_i + \delta_{\text{country}} + \delta_{\text{sector}} + \delta_{\text{year}}\big),$$

where $\Lambda(\cdot)$ is the logistic CDF. The control vector X_i includes:

- $\ln$(Offer Size), a standard IPO control capturing deal scale and placement difficulty;
- a missing indicator for offer size (and imputation to retain observations).

Fixed effects are included for country, industry sector, and year, capturing institutional differences across markets, systematic differences across sectors, and common macro regimes (e.g., crisis vs calm years). Standard errors are heteroskedasticity-robust (HC1). Because volatility is a market-level variable shared by IPOs occurring on the same day, date clustering is a natural robustness check; however, HC1 robust inference provides a transparent baseline and avoids very small-cluster issues that can arise when clustering by sparse event dates.

The coefficient β is interpreted in log-odds terms. To improve economic interpretability, the thesis also reports odds ratios $\exp(\beta)$ and predicted probabilities under different volatility states.

3.3.4 Main results

	Vol proxy	N	coef(z_vol)	se(HC1)	p-value	odds_ratio
0	V2X	3217	0.343815	0.084096	0.000043	1.410317
1	rolling30_vol_ann_pct	3217	0.127431	0.099720	0.201289	1.135907
2	ewma94_vol_ann_pct	3217	0.100248	0.102447	0.327811	1.105445
3	garch11_t_vol30_ann_pct	3217	-0.069632	0.104980	0.507149	0.932737
4	gjr11_t_vol30_ann_pct	3217	-0.053519	0.096724	0.580049	0.947888
5	egarch11_t_vol30_ann_pct	3217	-0.029119	0.091269	0.749694	0.971301

```
Avg predicted P(disruption) at z_V2X=-1.0: 0.068
Avg predicted P(disruption) at z_V2X=+0.0: 0.086
Avg predicted P(disruption) at z_V2X=+1.0: 0.109
Avg predicted P(disruption) at z_V2X=+2.0: 0.137
```

Figure 3.1 reports the results for six alternative volatility proxies (each regression includes the same controls and fixed effects). The key finding is that V2X implied volatility is strongly associated with IPO disruptions, while the backward-looking historical measures and the GARCH-family forecasts are not statistically significant in this specification.

V2X (implied volatility)

For the implied volatility proxy (V2X), the estimated coefficient is:

- $\beta = 0.3438$, HC1 SE $= 0.0841$, $p = 0.000043$
- Odds ratio: $\exp(0.3438) = 1.4103$

This estimate implies that a one standard deviation increase in V2X multiplies the odds of an IPO disruption by approximately 1.41, i.e., a 41% increase in disruption odds, controlling for offer size and fixed effects. This is economically meaningful and aligns with the thesis hypothesis: when the option market prices higher future volatility, IPO attempts are more likely to be postponed or withdrawn.

Historical volatility benchmarks (rolling 30-day and EWMA)

The historical volatility proxies have positive coefficients but are not statistically significant:

- Rolling 30-day volatility: $\beta = 0.1274$, $p = 0.2013$
- EWMA(0.94): $\beta = 0.1002$, $p = 0.3278$

The positive signs are directionally consistent with the hypothesis, but the absence of statistical significance suggests that backward-looking realized-style measures, once year/country/sector heterogeneity is controlled for, do not robustly explain disruption probability at the decision date in this sample.

GARCH-family forecasts

For the three GARCH-family proxies, the coefficients are slightly negative and not statistically significant:

- GARCH(1,1)-t: $\beta = -0.0696$, $p = 0.5071$
- GJR-GARCH(1,1)-t: $\beta = -0.0535$, $p = 0.5800$
- EGARCH(1,1)-t: $\beta = -0.0291$, $p = 0.7497$

Given the p-values, the appropriate interpretation is not that volatility forecasts reduce disruption risk, but that these parametric volatility forecasts do not show evidence of a systematic association with disruptions in this empirical design. The contrast between V2X and model-based forecasts is substantively interesting and suggests that the corporate finance channel may be more closely aligned with option-implied forward uncertainty than with purely statistical volatility forecasts based on past returns.

3.3.5 Economic magnitude: predicted disruption probabilities

Odds ratios are helpful but still abstract. To translate the V2X coefficient into more intuitive magnitudes, predicted probabilities were computed by holding all issuer characteristics and fixed effects at their observed values and setting the standardized V2X value to specific levels.

Average predicted disruption probabilities are:

- at $z_{V2X} = -1$: 6.8%
- at $z_{V2X} = 0$: 8.6%
- at $z_{V2X} = +1$: 10.9%
- at $z_{V2X} = +2$: 13.7%

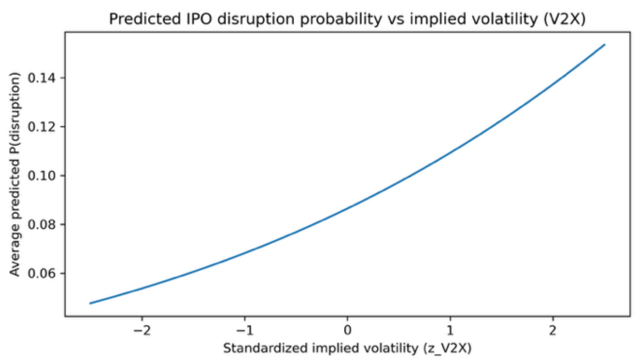

Figure 3.2: Predicted disruption probability vs V2X

These predictions imply that moving from an average volatility environment to a one-standard-deviation higher volatility environment raises the average disruption probability by about 2.3 percentage points (from 8.6% to 10.9%), which is a ~27% relative increase. Moving from a low-volatility state (-1σ) to a very high volatility state ($+2\sigma$) roughly doubles disruption risk (from 6.8% to 13.7%). This pattern is consistent with the financing-window view: as the market enters stressed, high-uncertainty regimes, the probability that IPO attempts fail to execute rises materially.

3.3.6 Interpretation

The fact that V2X is strongly significant while historical and GARCH-family measures are not has an intuitive interpretation that links directly back to Chapter 2.

First, implied volatility is forward-looking and embeds the market's aggregate pricing of uncertainty, including risk premia and investor demand for protection. The IPO decision to proceed or disrupt depends not only on the statistical expectation of volatility but also on broader risk appetite and distribution conditions—factors likely reflected in option prices. In this sense, V2X may measure a financing environment that is closer to the issuer's actual constraint set (investor willingness to allocate capital at a given price) than volatility estimates based solely on historical returns.

Second, Chapter 2's forecasting evaluation already suggested that implied volatility performs strongly as an out-of-sample predictor of subsequent volatility. The results here go further: they indicate that the same forward-looking measure is also corporate-finance relevant, predicting real financing disruptions at the deal level. This is exactly the "bridge" contribution of the thesis: volatility forecasting is not only a statistical exercise, but it has measurable consequences for firms' ability and willingness to access public equity markets.

Third, the lack of significance for backward-looking measures and GARCH forecasts may reflect a combination of factors: (i) these measures may be less salient in market practice compared with V2X, (ii) they may respond more slowly to sharp changes in risk sentiment, and (iii) once year fixed effects are included, much of the time-series variation they capture may be absorbed by macro regime controls. A natural next step (and a robustness check) is to test whether the results change under alternative timing (e.g., using filing date volatility) or alternative error structures (e.g., clustering by decision date).

Finally, it is important to note that disruption models capture only the extensive margin. Even when uncertainty rises, many deals still proceed to completion. The next section therefore studies pricing outcomes conditional on Trading, offer price revisions relative to the filing range and first-day returns (underpricing). Together, Sections 3.3 and 3.4 allow the thesis to distinguish whether volatility primarily operates through blocking issuance (disruptions) or through changing deal terms (pricing and underpricing) when issuance still occurs.

3.4 Empirical Models 2–3: Offer Price Revisions and First-Day Underpricing (Trading IPOs)

Sections 3.1–3.3 analysed the extensive margin of IPO execution: whether an IPO attempt is successfully completed or instead disrupted (withdrawn/postponed). The disruption results showed that option-implied volatility (V2X) is strongly associated with disruption risk, consistent with the thesis hypothesis that higher forward-looking uncertainty weakens the feasibility of raising equity at acceptable terms.

However, many IPOs proceed even in volatile periods. The next step is therefore to study the intensive margin, i.e., how volatility affects deal terms and immediate aftermarket performance conditional on completion. This section examines two outcomes that are central in IPO research and tightly connected to the financing frictions highlighted in Chapter 1: (i) offer price revisions during bookbuilding, and (ii) first-day underpricing (initial returns). Conceptually, these outcomes correspond to Schill's (2004) framing of offer price risk and aftermarket risk: when uncertainty is high, price discovery should be harder and underwriters may price more conservatively (or face larger aftermarket moves).

Importantly, these regressions are estimated only on Trading IPOs, meaning they are conditional on the deal not being disrupted. This conditionality is not a weakness; it is the correct object for pricing and underpricing. But it has an empirical implication: selection effects can attenuate (or even reverse) the unconditional relationship between volatility and pricing outcomes. If, during volatile states, only the strongest firms proceed to completion, the remaining sample may look "healthier" than the average IPO attempt in calm periods. This is one reason why it is valuable to analyse both (i) disruptions and (ii) conditional pricing outcomes: together they reveal where volatility matters most in the IPO process.

3.4.1 Sample and timing

The sample for this section includes IPO attempts with Offer Stage = Trading and with the relevant pricing variables available. The volatility proxy is measured at the pricing decision date, aligned exactly as in Section 3.2: the IPO is matched to the most recent trading day t on or before the pricing date, ensuring that volatility is observed contemporaneously and without look-ahead bias. The volatility measure used in the baseline specifications here is V2X implied volatility, standardized as a z-score:

$$zV2X_i = \frac{V2X_i - \overline{V2X}}{s(V2X)}.$$

Standardization is important because it allows the coefficient to be read as the effect of a one standard deviation increase in forward-looking market uncertainty.

3.4.2 Model 2: Offer price revision during bookbuilding

Definition of the dependent variable

Offer price revisions capture how the final offer price differs from the initial filing range communicated to investors. In practice, the filing range is updated during marketing, and the final offer price reflects bookbuilding information, demand conditions, and underwriter risk constraints.

Let P_i^{Offer} be the final offer price, and let P_i^{Low} and P_i^{High} be the initial filing range bounds ("Filing Term Price Range"). Define the midpoint as:

$$P_i^{Mid} = \frac{P_i^{Low} + P_i^{High}}{2}.$$

The offer price revision variable is typically operationalized as a percentage deviation from the midpoint:

$$Revision_i = \frac{P_i^{Offer} - P_i^{Mid}}{P_i^{Mid}},$$

which is positive when the final price is set above the midpoint (strong demand) and negative when set below (weak demand). The regression also includes the width of the filing range, which was computed as a percentage (or ratio) measure:

$$RangeWidth_i = \frac{P_i^{High} - P_i^{Low}}{P_i^{Mid}}.$$

Range width matters mechanically and economically. Mechanically, a wider range gives more "room" for price discovery, potentially reducing the need for large midpoint revisions. Economically, wider ranges may be chosen when valuation uncertainty is already high, allowing the issuer/underwriter to accommodate uncertain demand without large revisions.

Econometric specification

The baseline offer price revision regression (reported in Table 3.4) is estimated using OLS with heteroskedasticity-robust standard errors and includes fixed effects for country, sector, and year:

$$Revision_i = \alpha + \beta zV2X_i + \theta RangeWidth_i + \gamma \ln(OfferSize_i) + \delta' FE_i + \varepsilon_i.$$

Offer size enters because larger deals are typically more "institutional" and may face different price discovery and stabilization dynamics. Fixed effects control for structural differences across markets and sectors and for time-varying macro regimes.

term	coef	se	pval
Intercept	-0.1072829	0.038517132427510474	0.00534733
C(sector_grp)[T.Comm.]	0.040111586043289994	0.02066269	0.0522271
C(sector_grp)[T.Consumer, Cyclical]	0.01707544	0.019592126382667916	0.3834560794633213
C(sector_grp)[T.Consumer, Non-cyclical]	0.02476686	0.019273901164189654	0.1987941919802516
C(sector_grp)[T.Diversified]	0.00159521	0.020263682896196476	0.9372531613317561
C(sector_grp)[T.Energy]	0.011077526117604621	0.02114597	0.60037593
C(sector_grp)[T.Financial]	0.02843973	0.019540212806665624	0.14554576061368305
C(sector_grp)[T.Industrial]	0.022736933135598564	0.01977825	0.2503115366487647
C(sector_grp)[T.Technology]	0.040602337079951825	0.02037222	0.046258798368748424
C(sector_grp)[T.Utilities]	0.00524516	0.027515949732120633	0.8488213038900912
C(country_grp)[T.BD]	-0.0160227	0.03345278	0.6319640680087149
C(country_grp)[T.BE]	0.0443604	0.03182026	0.16328952726473578

C(country_grp)[T.CC]	0.03958865	0.032833601135254215	0.22791935798130059
C(country_grp)[T.CH]	0.03234702	0.031987800443670694	0.3119064858961642
C(country_grp)[T.DE]	0.07364957	0.03178783	0.0205086
C(country_grp)[T.F+]	0.015479542139740954	0.03230723	0.6318421935793523
C(country_grp)[T.FI]	0.04213358	0.03366611	0.21074736940720795
C(country_grp)[T.FR]	0.068371	0.02996263	0.022496686979312938
C(country_grp)[T.GB]	0.04067941	0.029643110027593393	0.16996829
C(country_grp)[T.GE]	0.03937036	0.02985557	0.18727143410527625
C(country_grp)[T.GR]	0.02121058	0.04538918	0.6402818711237499
C(country_grp)[T.GS]	-0.0042114	0.0377345	0.9111355919296299
C(country_grp)[T.IC]	0.09451965	0.03445423	0.00608178
C(country_grp)[T.IO]	0.04526453	0.040354202641132646	0.26199819875963515
C(country_grp)[T.IR]	0.027025691397797653	0.04403614	0.5394028785134374
C(country_grp)[T.IT]	0.03819324	0.029646745770944442	0.1976493491691993
C(country_grp)[T.JE]	-0.0201016	0.02968933	0.4983641644432971
C(country_grp)[T.LX]	-0.0055556	0.03901657	0.8867719464014929
C(country_grp)[T.MB]	0.11492485072583689	0.046829964740553884	0.014124244394876117
C(country_grp)[T.MQ]	-0.0932733	0.04485602	0.03758124
C(country_grp)[T.NE]	0.03766962	0.03306348	0.2545730002648413
C(country_grp)[T.NO]	0.022817029910709812	0.031211350072998426	0.4647491524071168
C(country_grp)[T.PO]	0.07493067	0.07562743	0.3217895999992483
C(country_grp)[T.RU]	0.17148780573951822	0.03190086	7.630749192097116e-08
C(country_grp)[T.SP]	0.026907873380431272	0.031231539835415212	0.3889291983799982
C(country_grp)[T.SW]	0.0612572	0.030192052215731346	0.04246664
C(country_grp)[T.SZ]	0.05676583	0.03090741	0.06626281
C(country_grp)[T.US]	0.021890686261231924	0.07524704	0.77111436
C(year_grp)[T.2011]	0.00484167	0.018757127717099312	0.7963109007084204
C(year_grp)[T.2012]	0.00680615	0.023280055281435615	0.7700116672997507
C(year_grp)[T.2013]	0.02868794	0.017572642724858097	0.10256696428901683
C(year_grp)[T.2014]	0.016568270450164235	0.017259812093822813	0.33708872501216536
C(year_grp)[T.2015]	0.01582774	0.016187480393898595	0.3281848250483568
C(year_grp)[T.2016]	0.00326591	0.017661080226204308	0.8532908299549906
C(year_grp)[T.2017]	0.028506997169076523	0.017684804166743777	0.10697310242220795
C(year_grp)[T.2018]	0.00920427	0.01690189	0.5860491038425002
C(year_grp)[T.2019]	0.0063884	0.018340123379622524	0.7275927899533546
C(year_grp)[T.2020]	0.016457411559819623	0.016976195272170206	0.3323254439384361
C(year_grp)[T.2021]	0.026637742991169024	0.016202068371089087	0.10015653736485707
C(year_grp)[T.2022]	-0.0046131	0.017484130102005346	0.7918995021206098

C(year_grp)[T.2023]	0.00252382	0.017650144333877748	0.8862968320361954
C(year_grp)[T.2024]	0.011546632886930131	0.01853809	0.5333764973437245
z_V2X	-0.0004884	0.00324559	0.8803917
range_width_pct	-0.3190537	0.031479740083585374	3.855554651586212e-24
ln_offer_size	0.00862458	0.00140473	8.269321275116397e-10
ln_offer_size_missing	-0.0176736	0.010462312023613607	0.0911694

Table 3.4 OLS regression output for offer price revisions

The key coefficient of interest is the one on standardized implied volatility:

- z_V2X: coefficient -0.0005, p-value 0.880 (95% CI: $[-0.007, 0.006]$)

This point estimate is essentially zero and statistically insignificant. In this baseline specification, there is no evidence that implied market volatility at pricing is associated with systematic offer price revisions relative to the filing range midpoint.

Two control variables, however, are strongly informative:

- Range width: coefficient -0.3191, p-value < 0.001

  This is economically intuitive: larger initial ranges are associated with smaller measured midpoint revisions. Interpreting magnitudes depends on the scaling of range_width_pct, but directionally the result indicates that a flexible initial valuation band reduces the need for subsequent repricing relative to the midpoint.

- Log offer size: coefficient $+0.0086$, p-value < 0.001

  Larger offerings are associated with slightly higher revisions (again depending on scaling, this is a modest effect). One interpretation is that larger, more visible issues may attract stronger institutional demand, supporting somewhat higher pricing relative to the initial midpoint.

Together with Section 3.3, the revision result suggests an important nuance: volatility is strongly associated with whether the IPO proceeds, but not necessarily with the within-process price revision conditional on completion. This can occur if underwriters and issuers self-select into completion only when they anticipate demand strong enough to price the deal without extreme revision, precisely the selection mechanism that makes the disruption model essential.

3.4.3 Model 3: First-day underpricing (initial returns)

Definition of the dependent variable

Underpricing is the classic IPO outcome: the return earned by investors on the first day of trading relative to the offer price. Using the first-day close price P_i^{Close}, underpricing is:

$$Underpricing_i = \frac{P_i^{Close} - P_i^{Offer}}{P_i^{Offer}}.$$

In the dataset this corresponds to Bloomberg "Offer to 1st Close" (or an equivalent field). Underpricing is interpreted as compensation to investors for participation and for bearing information and allocation frictions, and it is often higher under greater uncertainty.

Econometric specification

The baseline underpricing regression is:

$$Underpricing_i = \alpha + \beta zV2X_i + \gamma \ln(OfferSize_i) + \delta' FE_i + \varepsilon_i$$

estimated on Trading IPOs with robust standard errors and the same fixed effects structure (country/sector/year). In extended versions, the offer price revision as a control to separate demand revelation from aftermarket pricing effects could also be included.

term	coef	se	pval
Intercept	7.40570349	4.328452651924669	0.08709301
C(sector_grp)[T.Comm.]	-3.1936332	3.2475631330025005	0.3254137042408467
C(sector_grp)[T.Consumer, Cyclical]	-2.2862166	3.2536293526059166	0.4822635563581995
C(sector_grp)[T.Consumer, Non-cyclical]	-3.1525501	3.190881662287845	0.3231589371283685
C(sector_grp)[T.Diversified]	-6.2428657	3.8604006584577437	0.10584485544488345
C(sector_grp)[T.Energy]	-3.7847002	3.457646287656441	0.27369691753426517
C(sector_grp)[T.Financial]	-5.7233014	3.09469681	0.06440151
C(sector_grp)[T.Funds]	1.6725252051144137	4.464107800271104	0.7079129149129597
C(sector_grp)[T.Government]	-20.219731	8.101930168543603	0.012572005049631345
C(sector_grp)[T.Industrial]	-2.0862987	3.2382172228279704	0.5193978594355823
C(sector_grp)[T.Technology]	-0.2255349	3.329673501757063	0.9459967101735389
C(sector_grp)[T.Unknown]	3.885768813347464	7.22083073	0.5904850870411577
C(sector_grp)[T.Utilities]	3.64315741	7.026696629615222	0.6041278046322636
C(country_grp)[T.BD]	10.631139919640527	4.169281643949098	0.010776207586396034

C(country_grp)[T.BE]	-2.3908112	3.805294206270229	0.5298169203285942
C(country_grp)[T.CC]	0.8388349832431112	5.247396760566124	0.8729934397323782
C(country_grp)[T.CH]	1.3136439254374959	3.7103417907103884	0.7233019345592535
C(country_grp)[T.DE]	5.711415153295375	4.88620761	0.24244993698682904
C(country_grp)[T.F+]	1.2584034054919204	7.195132316951328	0.8611609663562189
C(country_grp)[T.FI]	2.516067713142542	3.961708214017387	0.5253653656364736
C(country_grp)[T.FR]	-0.9305075	3.355400496118074	0.7815371790015297
C(country_grp)[T.GB]	6.8571880111859445	3.268480979566864	0.035907444998973234
C(country_grp)[T.GE]	4.45596652	3.7253947547349213	0.2316552486432455
C(country_grp)[T.GI]	5.133703798536064	7.108277671251763	0.4701623854931979
C(country_grp)[T.GR]	4.783346805461213	7.51480253	0.5244354082324721
C(country_grp)[T.GS]	4.512089393826012	3.99652199	0.25889613746699625
C(country_grp)[T.IC]	5.9662126664889685	5.008175021204453	0.23353788975300638
C(country_grp)[T.IO]	25.03603551365531	15.388108981376185	0.10374286432747594
C(country_grp)[T.IR]	2.421498109333895	4.189726134193412	0.5632904817692993
C(country_grp)[T.IT]	4.67790281	3.354136607438296	0.16311635515299705
C(country_grp)[T.JE]	23.120014720153527	8.783259054429113	0.00848134
C(country_grp)[T.LC]	39.75671686551179	3.943696672687803	6.69865744333051e-24
C(country_grp)[T.LX]	-0.1779589	3.7174021668002286	0.9618183728831953
C(country_grp)[T.MB]	-6.3718416	5.198192669066181	0.22028140616405578
C(country_grp)[T.MQ]	-2.6444004	3.7914394202027197	0.4855111602863855
C(country_grp)[T.NE]	6.7008587107017705	4.5845654463358505	0.1438473943094995
C(country_grp)[T.NO]	1.6104949165880187	3.7154165457604282	0.6646785744900685
C(country_grp)[T.PO]	-0.9570618	3.705286119333782	0.7961782181852608
C(country_grp)[T.RU]	4.613464840316751	4.020662226498694	0.2512002103762293
C(country_grp)[T.SP]	2.8892812384662023	3.605105422233335	0.4228760754208484
C(country_grp)[T.SW]	1.194525271811758	3.510823950376842	0.7336752546227441
C(country_grp)[T.SZ]	9.071459813192513	4.163920386756423	0.029362416541031507
C(country_grp)[T.US]	19.141909787330675	4.538444352715727	2.46779628253125e-05
C(country_grp)[T.Unknown]	3.8857688133474957	7.220830731029452	0.5904850870411507
C(country_grp)[T.VS]	-24.399138	4.5043621188790555	6.06819574916769e-08
C(year_grp)[T.2011]	2.412575001499039	2.7805266556590564	0.38557591195016017
C(year_grp)[T.2012]	-1.8231747	2.1678846	0.4003521303533323
C(year_grp)[T.2013]	7.303373198816205	2.5823409909707866	0.00468108
C(year_grp)[T.2014]	1.5221423201151216	1.7235545120333098	0.37715988596819416
C(year_grp)[T.2015]	5.197760758642368	1.823769241120942	0.00437178
C(year_grp)[T.2016]	2.8731064967762476	2.0694119381830194	0.1650248414117751
C(year_grp)[T.2017]	4.664409606912977	2.2114364476839783	0.034925455020999654

C(year_grp)[T.2018]	1.0663764329788046	1.9898567429146583	0.5920234253718722
C(year_grp)[T.2019]	6.964380857150757	2.700232960627644	0.00990358
C(year_grp)[T.2020]	10.69208583928713	2.6443642746517186	5.26932501048239e-05
C(year_grp)[T.2021]	7.817027765706584	1.9597849204160118	6.64315685912935e-05
C(year_grp)[T.2022]	3.559168080157718	3.1679562142991693	0.26122931760967416
C(year_grp)[T.2023]	5.213393592298576	2.68539476	0.052211213514994456
C(year_grp)[T.2024]	8.966744619818371	3.6888427909057775	0.015066592576044879
z_V2X	0.19052774821258375	0.7063465180613538	0.7873626396704468
ln_offer_size	-1.269907	0.2929821406017759	1.46146384160692e-05
ln_offer_size_missing	5.987752341179184	10.321637250991468	0.5618360402209488

Table 3.5 OLS regression results for first-day underpricing

The key coefficient of interest is again the implied volatility term:

- z_V2X: coefficient +0.1905, p-value 0.787

This estimate is statistically insignificant and economically small relative to its standard error. Thus, conditional on completion, the baseline evidence does not show that higher implied volatility at the time of pricing translates into systematically higher (or lower) underpricing in first-day returns.

A strong and significant relationship appears for deal size:

- $\ln(OfferSize)$: coefficient -1.2699, p-value < 0.001

This negative association is common in IPO data: larger deals tend to have lower underpricing, consistent with lower information asymmetry, greater institutional participation, or stronger underwriter stabilization capacity. (The magnitude here is interpreted in the unit of the underpricing variable; if the underpricing is in percentage points, then the effect is in pp.)

Interpreting the joint pattern across Models 1–3

At a high level, the findings across the two parts of Chapter 3 are consistent with the following narrative:

- Volatility (especially implied volatility) is highly relevant for whether the IPO market is "open" for firms (Section 3.3: disruption probability rises materially with V2X).
- Once an IPO clears the hurdle to completion, pricing adjustments and first-day returns do not show a robust direct relationship with V2X in the baseline specification.

This is a coherent economic outcome: when volatility is elevated, many marginal deals do not price at all. The deals that do price are a selected subset for which demand is strong enough (or underwriters are confident enough) to proceed, which can reduce the sensitivity of revisions and underpricing to volatility.

4. Conclusions

4.1 Interpretation of Findings

This work asked a simple but underexplored question at the intersection of volatility forecasting and corporate finance: *does forward-looking market uncertainty, measured using volatility forecasts, shape real IPO financing outcomes in Western Europe?* The empirical analysis provides a clear answer, with an important nuance: *volatility matters most for whether an IPO can be executed at all,* while its role in explaining *conditional* pricing outcomes (revisions and first-day returns) is weaker once a deal reaches completion.

Volatility forecasting is not only statistically relevant, but economically relevant, when the right volatility is used.

A central empirical result is that *option-implied volatility contains the strongest corporate-finance signal* among the volatility measures considered. The disruption model shows that a one-standard-deviation increase in implied volatility is associated with a meaningful rise in the likelihood that an IPO attempt is *postponed or withdrawn* rather than completed.

In the baseline disruption specification, the implied volatility proxy has:

- a positive and highly significant coefficient ($\beta \approx 0.3438$, $p < 0.001$),
- an odds ratio of $\exp(0.3438) \approx 1.41$, implying that the odds of disruption increase by roughly 41% for a one standard deviation increase in implied volatility, all else equal.

Beyond statistical significance, the economic magnitude is transparent when translated into predicted probabilities. Average predicted disruption probabilities move from:

- 6.8% at a relatively calm volatility state (-1σ),
- to 8.6% at an average state (0σ),
- to 10.9% at $+1\sigma$,
- and 13.7% at $+2\sigma$.

This implies that moving from average conditions to a $+1\sigma$ volatility environment raises disruption probability by about 2.3 percentage points (8.6% → 10.9%), which is a substantial relative increase. Moving from a calm to a stressed state ($-1\sigma \rightarrow +2\sigma$) roughly doubles disruption risk (6.8% → 13.7%). These are economically meaningful changes in the success probability of accessing public equity markets.

Extensive-margin effects dominate: volatility predicts "execution risk," not "how the survivors are priced"

A second key finding is that the strong volatility effect appears primarily in Model 1 (disruption), not in Model 2 (offer price revisions) or Model 3 (underpricing) once we condition on completion (i.e., Trading IPOs only).

This pattern is informative. It suggests that volatility affects IPO financing mainly through the extensive margin: whether firms proceed, delay, or abandon issuance when uncertainty spikes. In contrast, conditional on proceeding to pricing and trading, underwriters and issuers may adapt pricing mechanisms (range setting, allocation, stabilization, and investor targeting) in ways that dampen any direct statistical relationship between contemporaneous volatility and realized pricing outcomes.

This is consistent with an IPO process view:

1. When uncertainty becomes too high, a non-trivial fraction of deals simply do not clear—hence the strong disruption effect.
2. Among deals that clear, pricing is already "filtered" through bookbuilding, investor feedback, and underwriter discretion; therefore the remaining cross-sectional variation in revisions/underpricing may be explained more by deal- and issuer-specific factors than by a single market-wide volatility indicator at pricing.

The disruption result is strongest for implied volatility, weaker for purely statistical volatility measures

The disruption model further shows that *historical volatility estimates and GARCH-family forecasts do not display robust explanatory power* in the same specification (once fixed effects and controls are included). Their coefficients are not statistically distinguishable from zero.

This contrast between implied volatility and return-based volatility proxies is substantively important. It suggests that the IPO disruption decision is tied not only to "expected realized volatility" in a narrow statistical sense, but also to broader market conditions that are closely related to:

- risk appetite,
- hedging demand,
- the market price of uncertainty,
- and distribution risk faced by underwriters.

Implied volatility plausibly embeds these dimensions more directly than backward-looking or purely parametric time-series estimates. In corporate finance terms, implied volatility appears to proxy the *financing environment*—not just the mechanical variability of returns.

Offer price revisions: no systematic volatility effect, but range-setting matters strongly

For offer price revisions relative to the filing midpoint, the estimated coefficient on implied volatility is essentially zero:

- z_V2X ≈ -0.0005, $p \approx 0.88$, with a confidence interval tightly centred around zero.

This does not support the idea that higher implied volatility at pricing mechanically forces larger midpoint revisions once the deal reaches that stage. Instead, the revision regressions highlight the importance of *range width* and *deal size*:

- Range width enters with a strong negative coefficient (≈ -0.319, $p < 0.001$), implying that wider filing ranges are associated with smaller midpoint revisions.
- Log offer size enters positively ($\approx +0.0086$, $p < 0.001$), suggesting that larger deals experience slightly higher revisions (interpretable as stronger demand revelation or greater placement capacity), though the economic magnitude depends on scaling.

The range-width result is particularly intuitive: when the issuer/underwriter chooses a more flexible valuation band at filing, the midpoint becomes a less binding reference point, so measured "revision

from midpoint" mechanically shrinks. Economically, this supports a view of the filing range as an instrument for managing valuation uncertainty and reducing the need for discrete repricing.
Underpricing: no volatility effect conditional on completion, but deal size and distributional features matter
For first-day underpricing (offer-to-first-close), implied volatility again does not show a reliable association:

- z_V2X ≈ +0.1905, p ≈ 0.79.

In contrast, deal size is strongly informative:

- ln(OfferSize) ≈ −1.27, $p < 0.001$, consistent with a common IPO regularity: larger, more institutionally placed offerings tend to have lower underpricing, potentially reflecting better information environments, more analyst attention, more stable investor bases, and/or underwriter capacity to price closer to fundamental demand.

A crucial descriptive finding is that first-day returns are highly dispersed in the sample:

- mean underpricing is around 10.2%,
- median around 3.1%,
- with extreme outliers (including very large positive values) and a lower bound reaching −100%.

This distributional shape matters for interpretation: it implies that underpricing outcomes are influenced by a minority of extreme observations, making it important to evaluate robustness to outliers (e.g., winsorization, trimming, robust regression, or quantile regression). The lack of volatility significance in a baseline OLS framework should therefore be interpreted as: there is no strong average linear relationship, not that uncertainty is irrelevant for every deal.

Taken together, the empirical results support a coherent narrative:

1. *Market uncertainty (especially as priced in options markets) significantly increases IPO execution risk.*
   Volatility primarily affects the probability that the firm can raise public equity at all, consistent with "financing windows" and underwriting/distribution risk channels.
2. *Conditional on completion, pricing outcomes are not mechanically tied to contemporaneous implied volatility.*
   This can arise because the bookbuilding mechanism and issuer-underwriter discretion manage uncertainty endogenously, and because the sample of completed IPOs is selected (especially in volatile periods).
3. *Institutional design choices (e.g., range width) and deal characteristics (e.g., offer size) explain more of the variation in revisions and underpricing than volatility at pricing.*
   This emphasizes the importance of separating market-wide conditions that influence *whether* a deal proceeds from mechanisms that determine *how* it is priced once it proceeds.

4.2 Implications for Corporate Finance

The findings have several implications for how firms and intermediaries should think about IPO timing, execution risk, and the role of volatility forecasts as decision tools.

Implications for issuers: volatility forecasts are most valuable as “go/no-go” and timing signals

The strongest empirical evidence relates to the disruption decision. For issuers, this suggests that volatility forecasts, especially forward-looking measures, should be treated as *execution-risk indicators*, not just background market descriptors.

Practically, this supports a decision framework where firms:

- monitor forward-looking uncertainty in the run-up to critical IPO milestones,
- update go/no-go decisions dynamically as market uncertainty changes,
- and explicitly evaluate the trade-off between delaying issuance (option value of waiting) versus the cost of missing a financing window.

The predicted probabilities offer a way to quantify the shift in execution risk across market regimes. Even if a firm does not “forecast” volatility formally, tracking a forward-looking market uncertainty measure can provide a measurable signal about the probability of successfully completing issuance within a given time window.

Implications for underwriters: distribution risk appears central and is linked to option-implied uncertainty

Underwriters face the operational consequences of volatile markets through:

- weaker and more dispersed demand curves during bookbuilding,
- higher inventory and stabilization risk,
- and greater uncertainty around aftermarket trading.

The result that implied volatility predicts disruptions aligns with the idea that option markets aggregate information about risk appetite and hedging demand that is directly relevant for distribution conditions. In practice, underwriters may therefore:

- adjust syndicate structure,
- alter marketing intensity,
- recommend wider filing ranges,
- or steer issuers toward postponement when implied uncertainty indicates poor distribution conditions.

The range-width finding supports the view that underwriters can use filing terms strategically to manage uncertainty and reduce the need for visible repricing relative to midpoint.

Implications for investors: volatility affects deal availability more than average first-day pricing outcomes

From the investor side, the evidence indicates that volatility spikes are more likely to reduce the supply of IPOs (via postponements/withdrawals) than to systematically increase average underpricing among completed deals. This suggests that volatility regimes may affect investors primarily through:

- the set of deals that reach the market,
- and potentially the composition/quality of deals (selection effects),

  rather than through a stable mechanical relationship with first-day returns.

Broader implication: integrating volatility forecasting and corporate finance improves “real-time” market understanding

A conceptual contribution of the thesis is to connect volatility forecasting, which is often evaluated only by statistical loss functions, to real corporate financing outcomes. The disruption evidence indicates that forward-looking volatility measures can serve as *real-time proxies for financing frictions*, not merely predictors of return variability.

This contributes to bridging two strands of literature:

- volatility modelling (where implied volatility is often strong empirically),
- and IPO market timing (where "market conditions" are central but frequently measured in coarse or backward-looking ways).

4.3 Limitations

While the findings are coherent and economically meaningful, several limitations should frame interpretation and motivate future extensions.

Measurement limitations in disruption timing

Disruption analysis relies on the best available disruption timing, but disruption dates can be incomplete for some deals. Missing disruption timing can reduce the usable disruption sample and may introduce selection if missingness is not random (for example, if smaller deals or certain markets have less reliable documentation). This limitation is important because the disruption model is central to the thesis.

A natural extension is to evaluate sensitivity to alternative timing definitions (e.g., using filing date or last update date) and to test whether results persist under different event-date assumptions.

Selection effects in pricing and underpricing regressions

The pricing and underpricing regressions are estimated on completed IPOs by construction. This is correct for the economic question conditional on completion, but it implies selection: in volatile states, weaker deals may drop out, leaving a set of stronger deals that are able to price and trade. Selection can attenuate (or sometimes reverse) unconditional relationships.

Future work could address this explicitly using:

- selection models (e.g., Heckman-style correction),
- hazard models of completion timing,
- or joint modelling of disruption and pricing outcomes.

Outliers and distributional non-normality in first-day returns

The underpricing distribution is heavy-tailed, with extreme observations. This motivates robustness checks such as:

- winsorizing first-day returns,
- trimming extreme percentiles,
- robust regression,
- or quantile regression to evaluate whether volatility matters in the tails even if the mean relationship is weak.

Omitted variables and richer issuer-level controls

While fixed effects and standard controls capture major cross-sectional differences, IPO outcomes are also influenced by issuer fundamentals, governance, underwriter reputation in more granular form, ownership structure, market sentiment proxies, and sector-specific cycles. Adding richer issuer-level and demand-side variables could refine inference about the volatility channel and help separate volatility from correlated market conditions.

4.4 Closing Remarks

This work set out to evaluate whether volatility forecasts, particularly forward-looking measures, shape IPO financing outcomes in Western Europe. The findings support a clear conclusion:

- *Forward-looking market uncertainty is strongly linked to IPO execution risk*
  When implied volatility rises, IPO attempts are materially more likely to be postponed or withdrawn. This is consistent with a financing-window view where volatility amplifies offer price risk and distribution risk to the point that issuers rationally delay or abandon issuance.
- *Conditional on completion, volatility plays a limited role in explaining average offer price revisions and first-day underpricing*
  Instead, institutional design choices (notably filing range width) and deal characteristics (notably offer size) are more informative for these outcomes. This pattern is consistent with endogenous adaptation and selection: volatile markets reduce the set of deals that reach completion, and the deals that do complete are priced through mechanisms that absorb uncertainty.

The broader implication is that volatility forecasting should not be viewed solely as a statistical forecasting exercise. In IPO markets, forward-looking volatility measures appear to operate as *state variables of the financing environment*—variables that shape the feasibility of external equity issuance in real time. This connection between volatility forecasting and corporate financing decisions is precisely where the thesis contributes: it shows that volatility information, especially when derived from market expectations, has tangible consequences for firms' access to public equity markets.

Bibliography

1 Allen, F., & Faulhaber, G. R. (1989). Signalling by underpricing in the IPO market. *Journal of Financial Economics, 23*(2), 303–323.

2 Andersen, T. G., & Bollerslev, T. (1997). *Yes, ARCH models do provide good volatility forecasts* (NBER Working Paper). National Bureau of Economic Research.

3 Andersen, T. G., Bollerslev, T., Diebold, F. X., & Ebens, H. (2001). The distribution of realized stock return volatility. *Journal of Financial Economics*.

4 Andersen, T. G., Bollerslev, T., Diebold, F. X., & Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*.

5 Archakov, I., Hansen, P. R., & Lunde, A. (2026). *A multivariate realized GARCH model*. Journal of Econometrics, 254.

6 Baker, M., & Wurgler, J. (2002). Market timing and capital structure. *The Journal of Finance, 57*(1), 1–32.

7 Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies, 16*(1), 101–143.

8 Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society: Series B*.

9 Barndorff-Nielsen, O. E., & Shephard, N. (2008). Designing realized kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica*.

10 Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., & Shephard, N. (2008). Realised kernels in practice: Trades and quotes. *Econometrics Journal*.

11 Baron, D. P. (1982). A model of the demand for investment bank advising and distribution services for new issues. *Journal of Finance, 37*, 955–976.

12 Beatty, R. P., & Ritter, J. R. (1986). Investment banking, reputation, and the underpricing of initial public offerings. *Journal of Financial Economics.*

13 Bekaert, G., Hoerova, M., & Lo Duca, M. (2013). Risk, uncertainty and monetary policy. *Journal of Monetary Economics, 60*(7), 771–788.

14 Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics, 98*(1), 85–106.

15 Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.

16 Blair, B. J., Poon, S.-H., & Taylor, S. J. (2001). Forecasting S&P 100 volatility: The incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*.

17 Boeh, K. K., & Dunbar, C. G. (2013). Post IPO withdrawal outcomes.

18 Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327.

19 Bollerslev, T., Chou, R. Y., & Kroner, K. F. (1992). ARCH modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics*.

20 Bollerslev, T., & Mikkelsen, H. O. (1996). Modeling and pricing long memory in stock market volatility. *Journal of Econometrics, 73*, 151–184.

21 Bollerslev, T., Tauchen, G., & Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies, 22*(11), 4463–4492.
22 Booth, J. R., & Smith, R. L. (1986). Capital raising, underwriting and the certification hypothesis. *Journal of Financial Economics, 15*, 261–281.
23 Brau, J. C., & Fawcett, S. E. (2006). Initial public offerings: An analysis of theory and practice. *The Journal of Finance, 61*(1), 399–436.
24 Brau, J. C., Ryan, P. A., & DeGraw, I. (2006). Initial public offerings: CFO perceptions. *Financial Review, 41*(4), 483–511.
25 Braun, P. A., Nelson, D. B., & Sunier, A. M. (1995). Good news, bad news, volatility, and betas. *The Journal of Finance, 50*, 1575–1603.
26 Busaba, W. Y. (2006). Bookbuilding, the option to withdraw, and the timing of IPOs. *Journal of Corporate Finance, 12*(2), 159–186.
27 Busaba, W. Y., & Benveniste, L. M. (1997). Bookbuilding versus fixed price: An analysis of competing strategies for marketing IPOs. Journal of Financial and Quantitative Analysis, 32(4), 383–403.
28 Busaba, W. Y., Benveniste, L. M., & Guo, R. (2001). The option to withdraw IPOs during the premarket: Empirical analysis. *Journal of Financial Economics, 60*, 73–102.
29 Canina, L., & Figlewski, S. (1993). The information content of implied volatility. *Journal of Finance.*
30 Carr, P., & Wu, L. (2006). A tale of two indices. *Journal of Derivatives, 13*(3), 13–29.
31 Chambers, D., & Dimson, E. (2009). IPO underpricing over the very long run. *Journal of Finance, 64*(3), 1407–1443.
32 Chemmanur, T. J., & Fulghieri, P. (1999). A theory of the going-public decision. *The Review of Financial Studies, 12*(2), 249–279.
33 Christensen, B. J., & Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of Financial Economics.*
34 Clark, T. E., & West, K. D. (2007). Approximately unbiased tests for nested model comparisons. *Journal of Econometrics.*
35 Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance, 1*(2), 223–236.
36 Corwin, S. A., & Schultz, P. (2005). The role of IPO underwriting syndicates: Pricing, information production, and underwriter competition. *Journal of Finance*, 60(1), 443–486.
37 Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics.*
38 Day, T. E., & Lewis, C. M. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics.*
39 de Souza e Almeida, V., & Leal, R. P. C. (2015). Brazilian initial public offerings, underwriters, and premium corporate governance segments listing. *Corporate Ownership & Control, 13*(1), 1410–1418.
40 Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics.*
41 Dixit, A., & Pindyck, R. (1994). *Investment under uncertainty.* Princeton University Press.
42 Drake, P., & Vetsuypens, M. L. (1993). IPO underpricing and insurance against legal liability. *Financial Management, 22*, 64–73.

43 Dunbar, C. G., & Foerster, S. R. (2008). Second time lucky? Withdrawn IPOs that return to the market. *Journal of Financial Economics.*
44 Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, *50*(4), 987–1007.
45 French, K., Schwert, G. W., & Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics, 19*, 3–29.
46 Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance,* 48, 1779-1801.
47 Grinblatt, M., & Hwang, C. Y. (1989). Signalling and the pricing of new issues. *Journal of Finance, 44*, 393–420.
48 Granger, C., & Poon, S.-H. (2005). Practical issues in forecasting volatility. Financial Analysts Journal, 61(1).
49 Hanley, K. (1993). The underpricing of initial public offerings and the partial adjustment phenomenon. *Journal of Financial Economics, 34*, 213–250.
50 Hansen, P. R., & Huang, Z. (2017). Exponential GARCH modeling with realized measures. *Journal of Business and Economic Statistics*.
51 Hansen, P. R., Huang, Z., & Shek, H. H. (2012). Realized GARCH: A joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*.
52 Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The model confidence set. *Econometrica* , *79*(2), 453-497.
53 Helbing, P., Lucey, B. M., & Vigne, S. (2018). The determinants of IPO withdrawal: Evidence from Europe. Journal of Corporate Finance.
54 Helwege, J., & Liang, N. (2004). Initial public offerings in hot and cold markets. *Journal of Financial and Quantitative Analysis, 39*(3), 541–569.
55 Ibbotson, R. (1975). Price performance of common stock new issues. *Journal of Financial Economics, 2*, 235–272.
56 Ibbotson, R. G., & Jaffe, J. F. (1975). "Hot issue" markets. *The Journal of Finance, 30*(4), 1027–1042.
57 Jiang, G. J., & Tian, Y. S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*.
58 Kambouroudis, D. S., McMillan, D. G., & Tsakou, K. (2016). Forecasting stock return volatility: A comparison of GARCH, implied volatility, and realized volatility models. *Journal of Futures Markets*, 36(12), 1127–1163.
59 Loughran, T., & Ritter, J. R. (1995). The new issue puzzle. *The Journal of Finance, 50*, 23–51.
60 Loughran, T., Ritter, J. R., & Rydqvist, K. (1994). Initial public offerings: International insights. *Pacific-Basin Finance Journal, 2*, 165–199.
61 Lowry, M. B. (2003). Why does IPO volume fluctuate so much? *Journal of Financial Economics, 67*, 3–40.
62 Lunde, A., & Hansen, P. R. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? Journal of Applied Econometrics, 20(7), 873–889.
63 Mandelker, G., & Raviv, A. (1977). Investment banking: An economic analysis of optimal underwriting contracts. *Journal of Finance, 32*(3), 683–694.
64 Mikkelson, W. H., & Partch, M. M. (1988). Withdrawn security offerings. *Journal of Financial and Quantitative Analysis, 23*(2), 119–133.

65 Modigliani, F., & Miller, M. H. (1958). The cost of capital, corporation finance, and the theory of investment. *American Economic Review, 48*, 261–297.
66 Myers, S. C., & Majluf, N. S. (1984). *Corporate financing and investment decisions when firms have information that investors do not have* (NBER Working Paper No. 1396). National Bureau of Economic Research.
67 Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* , *59*(2), 347–370.
68 Pagano, M., Panetta, F., & Zingales, L. (1998). Why do companies go public? An empirical analysis. *The Journal of Finance, 53*(1), 27–64.
69 Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business, 53*(1), 61–65.
70 Pástor, Ľ., & Veronesi, P. (2005). Rational IPO waves. *Journal of Finance, 60*(4), 1713–1757.
71 Pástor, Ľ., & Veronesi, P. (2009). Technological revolutions and stock prices. *American Economic Review, 99*(4), 1451–1483.
72 Patton, A. J., & Sheppard, K. (2009). Evaluating volatility and correlation forecasts. In T. G. Andersen, R. A. Davis, J.-P. Kreiss, & T. Mikosch (Eds.), *Handbook of financial time series.* Springer.
73 Poon, S.-H., & Granger, C. W. J. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature, 41*(2), 478–539.
74 Ritter, J. R. (1991). The long-run performance of initial public offerings. *The Journal of Finance, 46*, 3–27.
75 Ritter, J. R., Signori, A., & Vismara, S. (2013). Economies of scope and IPO activity in Europe. In M. Levis & S. Vismara (Eds.), Handbook of research on IPOs (pp. 11–34). Edward Elgar Publishing.
76 Ritter, J. R., & Welch, I. (2002). *A review of IPO activity, pricing and allocations* (Yale ICF Working Paper No. 02-01). Yale International Center for Finance.
77 Rock, K. (1986). Why new issues are underpriced. *Journal of Financial Economics, 15*, 187–212.
78 Schill, M. J. (2004). Sailing in Rough Water: Market Volatility and Corporate Finance. Journal of Corporate Finance, 10(5), 659–681.
79 Sherman, A. (1999). *Global trends in IPO methods: Book building vs. auctions* (Unpublished working paper). University of Notre Dame.
80 Tinic, S. M. (1988). Anatomy of initial public offerings of common stock. *Journal of Finance, 43*, 789–822.
81 Welch, I. (1989). Seasoned offerings, imitation cost, and the underpricing of initial public offerings. *Journal of Finance, 44*, 421–448.
82 Whaley, R. E. (2000). The investor fear gauge. *Journal of Portfolio Management, 26*(3), 12–17.
83 Wolfe, G. A., Cooperman, E. S., & Ferris, S. P. (1994). An analysis of the underwriter selection process for initial public offerings. *Journal of Financial Research, 17*(1), 77–90.