# Predictable Money, Unpredictable Returns?

## The Effects of US Federal Contracts on Equity Returns

**Ava Booker**
**Elna Jannati**

STOCKHOLM SCHOOL OF ECONOMICS
Department of Finance
4350 - Master's Thesis in Finance
Spring 2026

**Abstract:**

This thesis examines whether and how exposure to U.S. federal government procurement affects the equity returns of publicly traded firms. Drawing on contract obligations data from USASpending.gov matched to firm-level financial data from Compustat and CRSP, we construct a novel panel dataset covering 118 publicly traded firms across four major federal agencies, including the Department of Defense, the Department of Health and Human Services, National Aeronautics and Space Administration, and the Department of Veterans Affairs, over the period 2014 to 2024. Using a two-way fixed effects panel regression framework, we find no significant relationship between procurement exposure and stock returns in the full sample. However, disaggregating by agency reveals meaningful heterogeneity: NASA procurement exposure is associated with a statistically significant return discount, consistent with the interpretation that the market prices the superior cash flow stability of NASA's long-term, cost-plus contracting structure as lower risk. No equivalent effect is found for DoD, HHS, or VA. These findings suggest that treating federal procurement as a homogeneous exposure category conceals important cross-agency differences in how government revenue is priced by the market.

**Keywords:**
Government Procurement, Federal Contracts, Equity Returns, Asset Pricing, Cash Flow Stability, Agency Heterogeneity, United States

**Authors**:
Ava Booker (42500), Elna Jannati (42885)

**Supervisor:**
Laszlo Sajtos, Konjunkturinstitutet, House of Sustainable Society at the Stockholm School of Economics

**Examiner:**
Anders Andersson

**Acknowledgements:** We would like to thank our supervisor Laszlo Sajtos for his tremendous support and guidance throughout our thesis.

Master Thesis
MSc Program in Finance
Stockholm School of Economics

# Contents

# 1. INTRODUCTION

The U.S. government is one of the largest purchasers of goods and services in the global economy, with annual spending on contracts exceeding hundreds of billions of dollars every year. Compared with corporate customers, the U.S. government provides stable and long-term demand though procurement contracts and exhibits a very low risk of default (Dhaliwal, 2016), allowing recipient firms to benefit from more stable revenues, and consequently more predictable cash flows (Goldman, 2020). For the firms on the receiving end, this raises a natural question: does it matter? If a firm derives a substantial share of its revenue from a customer that almost never defaults, operates on multi-year contractual commitments, and is backed by the full fiscal authority of the federal government, should we expect that to show up in how its stock is priced?

This is not merely an academic question. Government procurement constitutes a major channel through which public resources are allocated to private firms in the United States and may have a direct impact on the performance of firms' who receive these contracts (Goldman, 2020). Understanding whether and how markets price this exposure matters for investors assessing risk, for policymakers evaluating the broader economic effects of procurement decisions, and for anyone curious about whether the predictability of public demand translates into measurable financial advantages for the firms that depend on it. These considerations naturally raise the question of whether financial markets recognize and price the stability associated with government demand.

Within the framework of the Efficient Market Hypothesis (EMH), asset prices reflect all available information (Fama, 1970), while standard asset pricing models, such as the Capital Asset Pricing Model (CAPM), predict that expected returns are determined by exposure to market risk, typically interpreted as systematic risk (Sharpe, 1964). Given the stability of government demand, one may argue that firms subject to procurement are less sensitive to macroeconomic fluctuations. To the extent that this reduces their exposure to systematic risk, such firms would be expected to exhibit lower expected returns relative to firms whose revenues are more closely tied to cyclical demand.

A strand of literature supports this stabilizing view of government procurement. Firms with government contracts experienced smaller declines in performance during the 2008 - 2009 crisis (Goldman, 2020), while higher defense spending is associated with lower stock market volatility in exposed industries (Cortes el al., 2022). At the firm level, government customers are linked to a lower cost of equity and reduced risk relative to those with a higher corporate customer concentration (Dhaliwal et al., 2016). Taken together, these studies suggest that government procurement stabilizes firm performance and is associated with lower volatility and perceived risk. Under standard asset pricing theory, lower risk exposure should be associated with lower expected returns.

However, another strand of literature documents positive abnormal returns associated with government procurement. Event study evidence shows that contract award announcements generate short-term positive stock reactions (Diltz, 1990; Samunderu & Yordanova, 2024; Larson & Picou, 2002). Beyond announcement effects, firms with higher government sales or larger contract awards exhibit persistent outperformance not explained by standard risk factors (Parajuli, 2023; Pyun, 2025). Similarly, firms subject to government procurement earn returns excess of those predicted by standard asset pricing models over political cycles (Belo et al., 2013). This evidence is consistent with a potential mispricing mechanism, whereby the market does not fully incorporate the implications of government demand in asset prices.

Taken together, these findings reveal an unresolved tension in the literature. On the one hand, stable government demand should reduce firms' systematic risk, which under standard asset pricing theory should be associated with lower expected returns. In an efficient market, these risk characteristics should already be reflected in asset prices. On the other hand, empirical evidence points to positive abnormal returns for government exposed firms. This tension in the literature raises the question of whether government procurement primarily affects expected returns through changes in risk exposure or through predictable market underreaction to public demand shocks.

An existing gap in the literature is that government procurement is treated as a homogenous source of demand or focuses on a specific set of contracts or industries. In reality, federal agencies differ significantly in terms of their budget structure, contracting mechanism, demand drivers and exposure to macroeconomic and political shocks. For instance, The Department of Defense (DoD) is characterized by large, multiyear procurement (MYP) contracts (O'Rourke, 2025) and stable demand (Cortes et al., 2022). Whereas the National Aeronautics and Space Administration (NASA) exhibits greater project-level uncertainty (Lindbergh, 2026) and relies more heavily on discretionary funding (Dreier, 2026). This paper argues that the existing tension in the literature may arise from treating government procurement as homogenous, while in reality agency-level differences may generate heterogeneous risk exposures and, consequently, different stock returns.

This paper addresses the existing gap by introducing agency-level heterogeneity into the analysis of government procurement exposure. To do so, a panel dataset is constructed by combining contract level obligations data from USASpending.gov with firm level financial data from Compustat and stock return data from CRSP. The analysis focuses on publicly traded firms with exposure to four major agencies: The DoD, the Department of Health and Human Services (HHS), NASA and the Department of Veterans Affairs (VA). Procurement exposure is measured as the ratio of contract obligations to firm revenue, allowing for consistent comparison across firms and over time. The empirical framework employs a panel regression approach with firm year fixed effects to examine how variation in procurement exposure relates to stock returns.

The findings suggest that the relationship between government procurement and stock returns is not uniform across agencies. While no significant effect is detected in the full sample, NASA procurement exposure is associated with a statistically significant return discount. No equivalent effect is found for DoD, HHS, or VA. These results suggest that treating federal procurement as a homogenous category may cause meaningful differences to go undetected.

# 2. LITERATURE REVIEW

## 2.1 Government procurement exposure and asset pricing

The extent to which government procurement affects firm-level risk and stock returns has attracted considerable attention in the literature, yet evidence remains mixed. On the one hand, revenue stability associated with government contracts suggests that exposed firms should exhibit lower systematic and potentially lower idiosyncratic risk and, consequently, lower expected returns. On the other hand, several empirical studies document positive abnormal returns among firms subject to government procurement. The following section reviews both strands of evidence and shows how this tension motivates the present study.

## 2.2 Evidence suggesting lower risk and returns

Goldman (2020) studies U.S. publicly listed firms that derive more than 10% of their annual revenues from government procurement. Exploiting the 2008-2009 financial crisis as a sudden and plausibly exogenous shock, the paper shows that these firms experienced significantly smaller declines in sales, profitability, investment and employment relative to otherwise similar firms. In addition, firms with significant reliance on government demand maintained high market valuations during the crisis, with market capitalization approximately 15% higher than comparable firms by 2009. These findings suggest that stable government procurement acts as a hedge against macroeconomic shocks, highlighting the role of the government not only as a stable, but also a reliable source of demand.

Complementing the above, Cortes et al. (2022) study more than a century of U.S. military expenditure and show that higher defense spending is associated with lower aggregate U.S. stock market volatility. They attribute this effect to the stabilizing role of government demand, which reduces uncertainty regarding firms' expected cash flows and makes future earnings easier to forecast. The effect is strongest among industries closely tied to military production and procurement, while some related sectors also experience spillover effects.

Further evidence on the role of customer composition is provided by Dhaliwal et al. (2016), who examine U.S. publicly listed industrial firms over the period 1981 - 2011. Dhaliwal et al. (2016) demonstrate that high corporate customer concentration increases the cost of equity due to increased exposure to demand shocks and default risks, government customers have the opposite effect. The authors find that firms deriving at least 10% of revenue from the federal government see a 19 basis point reduction in the cost of equity, which Dhaliwal et al. (2016) attribute to the relative stability and longer duration of government contracts.

Taken together, these studies suggest that government procurement is associated with more stable cash flows and lower exposure to economic uncertainty, which may translate into potentially lower systematic and idiosyncratic risk. Under standard asset pricing theory, lower exposure to systematic risk should result in lower expected returns for firms subject to government procurement.

## 2.3 Evidence of anomalous returns and potential mispricing

A second strand of the literature documents positive abnormal returns for firms exposed to government procurement. Parajuli (2023) finds that firms for which government sales account for at least 10% of revenues earn significantly positive stock returns of 6% per year compared to firms without such exposure. These return differentials persist after controlling for a broad set of factors including policy uncertainty, geopolitical risk, and political partisanship cycles. Parajuli (2023) interprets the results as consistent with investor underreaction to publicly available information on government procurement. Underreaction implies a mispricing mechanism, whereby investors underreact to publicly available information on government procurement and fail to incorporate its implications for future profitability and cash flow stability into stock prices.

Building on this line of research, Pyun (2025) finds that firms receiving larger government contracts tend to earn higher stock returns over the six months following contract awards.. Portfolio evidence shows that firms receiving larger contract awards outperform both firms with smaller awards and the broader market. Pyun (2025) suggests that contract size is a key driver of the observed outperformance, particularly among large capitalization firms. The findings are consistent with the view that government contracts serve as a signal of firm strength and future growth prospects, contributing to higher expected cash flows and improved stock performance over time.

Belo et al. (2013) show that the performance of firms with different levels of exposure to government spending varies systematically over U.S. presidential cycles. Specifically, industries with high government exposure earn significantly higher stock returns during Democratic presidencies, while the pattern reverses under Republican administrations. An investment strategy that exploits these predictable differences in cross-sectional stock returns generates significant abnormal differences of 6.9% per year. These abnormal returns remain robust after controlling for firm level characteristics, business cycle fluctuations and standard asset pricing factors. Belo et al. (2013) interpret the findings as evidence consistent with gradual market underreaction to predictable changes in government spending policies.

While the above studies focus on long-term return predictability, a related strand examines short-term market reactions to government contract announcements. Samunderu & Yordanova (2024) examine NASDAQ-listed pharmaceutical companies and find that announcements of U.S federal government contracts lead to positive and statistically significant cumulative abnormal returns, although the effects are short-lived. Similarly, Larson & Picou (2002) document positive abnormal returns around contract announcements from military, non-military, municipal and foreign government, as well as corporate grantors. While abnormal returns are statistically significant for military contracts , they are concentrated around the announcement date and fully incorporated into prices within two trading days, with no evidence of post-announcement drift. Lastly, Diltz (1990) provides early evidence that U.S. DoD contract awards are associated with positive abnormal returns concentrated around the announcement date. However, when all contract types are pooled, the effects are not statistically significant, suggesting that the impact varies across contract characteristics.

## 2.4 Discussion of Literature

The evidence reviewed above reveals a tension in the asset pricing implications of government procurement exposure. On the one hand, government procurement is associated with more stable cash flows, lower cost of equity and reduced exposure to macroeconomic shocks (Dhaliwal et al., 2016; Cortes et al., 2022; Goldman, 2020), properties that under standard asset pricing theory should be associated with lower expected returns. On the other hand, empirical studies consistently document positive abnormal returns among firms exposed to government procurement, both at the moment of contract announcements (Diltz, 1990; Samunderu & Yordanova, 2024; Larson & Picou, 2002) and over longer horizons (Belo et al., 2013; Parajuli, 2023; Pyun, 2025), that are not fully explained by standard risk factors. Taken together, the existing findings suggest that while financial markets react quickly to specific contract announcements, they might not fully incorporate the broader and more persistent implications of government procurement exposure for a firm's future cash flows and risk profiles. An understudied aspect of existing literature is that it largely treats government procurement as a homogenous source of demand, which can serve as an explanation to the potentially conflicting findings. This suggests that the impact of government procurement on firms risk and return is complex and potentially heterogeneous, motivating a more nuanced analysis of how differences across government agencies shape firm's risk exposures and expected returns.

To develop this perspective, the following section examines how differences in budget structures, contracting mechanisms and demand drivers across major U.S. federal agencies may lead to different implications for firm risk and expected returns. By introducing agency-level heterogeneity, the analysis moves beyond viewing government procurement as a homogenous source of demand and aims to capture the differences across agencies shape firms' stock returns.

## 2.5 Agency Level Heterogeneity in Government Procurement

The DoD represents a case of large-scale procurement, with obligations approaching nearly $167.6 billion in fiscal year 2025 (Neenan, 2026). Its spending remains relatively stable, supported by long-term contracting mechanisms such as MYP contracts, allowing a single contract to span two to five years and providing a stable and predictable planning environment for contractors (O'Rourke, 2025). Defense spending has also demonstrated countercyclical properties, with Goldman (2020) showing that procurement acted as a hedge during economic downturns, and Cortes et al. (2022) documenting that higher military spending is associated with lower stock market volatility in defense-linked industries. While acquisition programs face execution risks such as cost overruns and delays (Government Accountability Office [GAO], 2025a), the structural stability of demand is expected to dominate, suggesting that DoD exposed firms should exhibit lower cash flow uncertainty and reduced sensitivity to market movements.

NASA procurement is project-based which implies procurement spending is closely linked to the development of individual programs or projects, rather than being based on pre-determined long-term funding targets and the agency relies on annual appropriations (Lindbergh, 2026). NASA's acquisition management has remained on GAO's High Risk List over three decades due to cost overruns and schedule delays across major programs (GAO, 2025b). Together, these features suggest that NASA procurement generates less predictable cash flows than DoD procurement, with greater exposure to program-level and budget uncertainty, implying higher firm specific risk for contractors.

The VA differs fundamentally, as procurement demand is driven primarily by demographic and policy demand. The number of veterans enrolled in and utilizing VA healthcare has remained relatively stable, driven by expanding eligibility rather than macroeconomic conditions (Panangala & Sussman, 2024; Eibner et al., 2016). This generates stable recurring demand, closely matching the government customer relationship that Dhaliwal et al. (2016) associate with lower cost of equity and reduced earnings uncertainty.

Lastly, the HHS exhibits a dual procurement structure that sets it apart from the other three agencies. On the one hand, some subprograms, such as Medicare and Medicaid are funded through mandatory spending and generate stable, recurring demand. On the other hand, emergency driven procurement is irregular and concentrated in crisis periods (Centers for Medicare & Medicaid Services, 2026; Sekar & Sheikh, 2025). This duality means that the risk implications of HHS exposure depend heavily on which procurement channel a firm is exposed to, and HHS exposed firms are therefore expected to show the widest dispersion and risk and return outcomes across the sample.

Taken together, these differences suggest that the asset pricing implications of government procurement are inherently heterogeneous. The stability and persistence of demand vary considerably across agencies, and these structural differences should be reflected in the cash flow uncertainty, systematic risk exposure and ultimately the expected returns of firms subject to procurement. Treating federal procurement as a homogeneous category risks averaging out variation that is important for understanding the conflicting findings in existing literature.

# 3. METHODOLOGY

## 3.1 Overview

This report documents the methodology used to construct an analysis-ready panel dataset for studying the relationship between government procurement exposure and firm-level equity returns. The central research question is whether, and to what extent, a firm's dependence on federal contract revenue affects its stock return, idiosyncratic return volatility, and systematic risk (market beta). To answer this, the study requires three things to be measured simultaneously for the same firm: how much government money the firm receives, how large the firm's total revenue is, and what its stock is doing. None of these come from the same source, and each source uses a different identifier for the same firm. The methodology described here is primarily the work of pulling all three together into a single, consistent panel.

The study covers four federal agencies, the Department of Defense (DoD), the Department of Health and Human Services (HHS), the National Aeronautics and Space Administration (NASA), and the Department of Veterans Affairs (VA), over the period 2014-2024. These agencies were selected because they collectively span the major categories of federal procurement: defense hardware and services, healthcare and pharmaceuticals, high-technology aerospace, and medical/IT services to veterans. Together they represent the majority of discretionary federal contract spending with publicly traded firms.

The construction process had five stages: (1) contract data extraction from the USASpending.gov API, (2) entity name standardization to consolidate subsidiary-level records to publicly traded parents, (3) firm eligibility filtering to retain only the publicly traded parent companies identified in the name standardization step, (4) financial data retrieval from WRDS, and (5) panel merging and exposure variable construction. Each stage is documented in full below, including the specific design decisions made and the rationale for each.

## 3.2 Data Sources

### 3.2.1 Contract Data: USASpending.gov

Contract award data was obtained from USASpending.gov, the U.S. federal government's official procurement disclosure platform, which is maintained by the Office of Management and Budget under the Federal Funding Accountability and Transparency Act. All federal agencies are legally required to report contract awards to this system, making it the authoritative source for procurement obligations data.

Data was retrieved via the platform's public REST API (version 2). The API accepts structured queries specifying an agency, a fiscal year range, and an award type filter, and returns a ranked list of the top contract recipients by total obligated dollar amount. The specific endpoint used was the recipient aggregation endpoint, which must be called programmatically, and it cannot be accessed by clicking a link directly in a browser. The endpoint is as follows:

```
POST https://api.usaspending.gov/api/v2/search/spending_by_category/recipient_duns/
```

This endpoint returns a ranked list of the top contract recipients by total obligated dollar amount for a given agency and time period. The term 'obligated amount' refers to the amount the government has formally committed to pay under the contract, it is legally binding and recorded at the time of award. This is distinct from amounts actually disbursed or invoiced, but it is the standard measure used in the procurement literature because it captures the full value of the government's commitment at the time it is made. While actual disbursements more closely track cash inflows, contract obligations are publicly disclosed at award and are therefore observable by investors at the time of commitment. The procurement literature accordingly uses obligations as the primary measure of exposure, on the grounds that the market can price the government's commitment when it is made rather than waiting for invoicing to occur.

Queries were restricted to definitive contracts only, using award type codes A (BPA Call[1]), B (Purchase Order), C (Delivery Order), and D (Definitive Contract). Definitive contracts are contracts where the scope of work, delivery schedule, and price are all fully agreed upon at the time of award, as opposed to open-ended vehicles whose eventual value is uncertain. This excludes grants, cooperative agreements, loans, and indefinite delivery vehicles that have not yet been called down, the latter would inflate apparent obligations with contract ceiling values that

[1] A BPA call is a call against a pre-existing Blanket Purchase Agreement, representing a firm order placed under a previously established framework contract.

may never be spent. The query covered fiscal years 2014 through 2024, with each U.S. federal fiscal year running from October 1 to September 30.

A note on the DUNS identifier: the API endpoint path includes the term 'recipient_duns', which refers to the Data Universal Numbering System, a nine-digit identifier historically used by the federal government to track contractors. This is an artifact of the API's naming convention; the actual response data contains recipient legal names and obligation amounts rather than DUNS numbers themselves. The federal government transitioned to a new identifier system (the Unique Entity Identifier, or UEI) in April 2022, but the API endpoint name was retained for backward compatibility. The data content returned is unaffected by this naming convention.

For each agency, up to 1,000 recipient records were retrieved by paginating through 10 pages of 100 results each. This captures all recipients with meaningful contract volume; the 1,000th recipient in the DoD results, for example, had cumulative obligations well below the threshold that would make a firm a significant procurement counterparty for the purposes of this study. The same held true for HHS, NASA, and VA, where the recipient lists were considerably shorter and the cap was never a binding constraint.

### 3.2.2 Connection Stability and Retry Logic

The USASpending API is a public, rate-limited service[2] that does not guarantee connection stability for large sequential requests. During development, connection drops were encountered consistently when running multi-page queries for multiple agencies in sequence. To ensure the integrity of the pull, the fetch script was designed with four specific resilience mechanisms.

First, a 30-second per-request timeout was set so that the script fails cleanly rather than hanging indefinitely on a dropped connection. Second, a retry loop with exponential backoff[3] was implemented: if any individual page request fails, the script waits 10 seconds and retries up to three times before moving on. Third, after every successfully retrieved page, the partial results are written to a checkpoint CSV file (e.g., dod_checkpoint.csv). If the script crashes at any point, it can be restarted and will detect the checkpoint file, load the rows already retrieved, and resume from the next page rather than starting over. Fourth, a 2-second delay was inserted between each successful page request to reduce the likelihood of the server interpreting the sequential requests as abusive traffic and closing the connection.

The specific parameter values, like the 30-second timeout, 10-second retry wait, three retry attempts, and 2-second inter-page delay, were determined empirically during development: initial runs without these safeguards produced consistent failures at pages 4-7 of the HHS and DoD

[2] A rate-limited service means the server restricts how many requests can be made within a given time window and will reject or drop connections if that threshold is exceeded. This limitation ensures that the API server remains stable.
[3] Exponential backoff is a strategy in which the waiting period between successive retry attempts increases with each failure, so that a script which fails once waits 10 seconds, and if it fails again it waits progressively longer before the next attempt, reducing pressure on a server that may be temporarily overloaded.

pulls, and the values above were the minimum settings that produced stable, complete retrieval across all four agencies.

These mechanisms are not merely defensive, they are necessary for reproducibility. Without them, a script that crashes at page 7 of the HHS pull would silently produce a truncated dataset with no indication that records were missing.

### 3.2.3 Stock Return Data: CRSP

Monthly stock return data was obtained from the Center for Research in Security Prices (CRSP) Monthly Stock File (crsp.msf) via the WRDS platform. CRSP is the standard source for equity return data in academic finance and provides total returns (inclusive of dividends), prices, shares outstanding, and trading volume for all securities listed on major U.S. exchanges. The query covered January 2013 through December 2024 for all 118 firms in the sample (identified by their CRSP permanent firm identifier, the PERMNO).

A common but important error in return-based studies is ignoring delisting returns, the return a shareholder receives when a stock is removed from trading due to acquisition, bankruptcy, or going-private. CRSP stores these separately in the delisting file (crsp.msedelist) because they often occur mid-month and cannot be recorded in the regular monthly file. If ignored, the last regular monthly return before delisting appears normal, understating the true loss for firms that delist at distressed prices. Delisting returns were retrieved and merged into the monthly file. Where a delisting return was missing, it was set to -1 (a total loss), which is a conservative but standard assumption. Where both a regular monthly return and a delisting return existed for the same month, they were compounded as (1 + ret) × (1 + dlret) - 1.

### 3.2.4 Financial Statement Data: Compustat

Annual financial statement data was obtained from Compustat via WRDS. Compustat is the standard source for firm fundamentals in academic finance and covers essentially all publicly traded firms with annual and quarterly income statements, balance sheet, and cash flow data. Two separate Compustat databases were queried: comp.funda (Compustat North America) for U.S.- and Canada-listed firms, and comp.g_funda (Compustat Global) for the eight foreign-listed ADRs[4] in the sample (BP, GSK, Sanofi, AstraZeneca, Philips, Roche, Novartis, and BAE Systems). The North America and Global databases use slightly different column names for equivalent variables, for example, shares outstanding is csho in North America and cshr in Global, and fiscal year-end price is prcc_f in North America and prc in Global. Both were queried separately and stacked into a single fundamentals table.

[4] American Depository Receipt. These are securities issued by a U.S. bank representing shares in a foreign company, allowing them to trade on U.S. exchanges.

The query was filtered to consolidate financial statements only (consol = ‘C’), which means the parent company’s financials inclusive of all subsidiaries, as opposed to the parent-only standalone statements. This is the correct filter for a study of publicly traded parents because it reflects the economic entity that shareholders actually own. Variables retrieved include net sales (the primary denominator for the exposure ratio), total assets, long-term debt, R&D expenditure, operating income, capital expenditure, cash, shares outstanding, and fiscal year-end stock price. The selection of control variables follows the standard specification in empirical asset pricing studies of firm-level returns and the procurement literature specifically (Goldman, 2020; Parajuli, 2023; Dhaliwal et al., 2016).

### 3.2.5 Identifier Bridge: CRSP-Compustat Merged Link Table

CRSP and Compustat use entirely different proprietary identifiers for the same firm, CRSP uses the PERMNO and Compustat uses the GVKEY. The CRSP Compustat Merged, or CCM, link table is the official bridge that resolves this problem, as it is a WRDS-maintained database that maps each Compustat GVKEY to its corresponding CRSP PERMNO, along with the date range over which that mapping is valid. In the CCM link table, each row records a GVKEY-PERMNO pair along with the date range over which the link is valid, allowing for firms that changed their primary listing or were involved in mergers. Link types LC (hand-verified) and LU (algorithmically verified) were retained, and only primary securities (linkprim = ‘P’ or ‘C’) were kept to avoid matching a firm’s GVKEY to a secondary class of shares or a preferred stock PERMNO.

## 3.3 Agency Selection and Rationale

The selection of agencies for this study was not arbitrary. Each agency was chosen because it represents a distinct economic mechanism through which government spending enters firm revenue, and because its contractor base contains a sufficient number of publicly traded firms to support meaningful statistical analysis. The following table summarizes the selection rationale.

| Agency | Procurement Budget | Mechanism | Primary Firm Types in Sample |
|---|---|---|---|
| Dept. of Defense (DoD) | ~$400B/yr | Long-term platform contracts (ships, aircraft, vehicles, IT systems) create multi-year committed revenue streams. Contract cancellations require Congressional action and are politically costly, giving high revenue predictability. | Aerospace & defense primes, shipbuilders, IT services, logistics, ammunition |

| | | | |
|---|---|---|---|
| Dept. of Health & Human Services (HHS) | ~$30-80B/yr (highly variable) | Medical countermeasure stockpiling, IT modernization, and managed care administration. COVID-era vaccine contracts created a large natural experiment in sudden exposure change. | Pharma/biotech, diagnostics, health IT, managed care, laboratory services |
| NASA | ~$20B/yr | High-technology development and operations contracts for space systems, launch vehicles, and aeronautics research. Contracts are concentrated among a small set of specialized aerospace firms with high revenue dependence. | Launch vehicles, satellites, aerospace engineering, space data services |
| Dept. of Veterans Affairs (VA) | ~$25B/yr | Medical equipment procurement, pharmaceutical supply, and IT modernization for the U.S. veterans healthcare system. Introduces a distinct set of healthcare and technology firms not fully captured by HHS. | Medical devices, pharmaceuticals, enterprise IT, health IT |

The Department of Energy (DoE) was explicitly considered and rejected. DoE’s procurement is dominated by the management and operation of government-owned, contractor-operated (GOCO) national laboratories, facilities such as Los Alamos, Lawrence Livermore, Sandia, and Oak Ridge that are owned by the federal government but managed by contractors. The entities that receive these management contracts are typically non-profit consortia (e.g., Battelle Memorial Institute, the University of California system), non-publicly-traded private firms (e.g., Bechtel National), or FFRDCs with no equity. Because the study requires matching contract recipients to publicly traded equity, DoE’s contractor base is structurally unsuitable.

A related filtering challenge arose in cases where a publicly traded firm participated in a consortium or joint venture alongside non-publicly-traded partners. In these arrangements, the total contract obligation is recorded under the lead recipient's legal name in USASpending, which may be the consortium itself rather than the publicly traded member. Where the lead recipient was identified as a non-publicly-traded consortium entity, the record was excluded from the sample on the grounds that the publicly traded partner's economic exposure to the contract could not be reliably isolated from the total obligated amount. Where the publicly traded firm was itself the named prime contractor, with non-public firms appearing as subcontractors, the full obligated amount was retained, consistent with the standard practice in the procurement literature of attributing prime contract value to the named awardee. This conservative approach means that in some cases a publicly traded firm's true government exposure may be marginally understated, since subcontract revenue flowing from consortium-led awards is not captured. This limitation is acknowledged but is unlikely to materially affect results given that the firms most central to the study, including defense primes, large pharma, and aerospace majors, are typically named as prime contractors rather than consortium members

The four selected agencies collectively represent the dominant channels of federal procurement with the private sector and span defense, healthcare, technology, and space, four distinct industries with different risk profiles, contract structures, and market dynamics. This cross-agency design is one of the study's primary contributions relative to the existing literature, which predominantly focuses on DoD alone.

## 3.4 Entity Name Standardization

### 3.4.1 The Name Mapping Problem

The most substantive data challenge in this study is that the federal government records contract obligations at the subsidiary level, meaning the specific legal entity that signed the contract, rather than at the level of the publicly traded parent. This creates a fundamental mismatch between the unit of observation in USASpending (a contracting subsidiary) and the unit of analysis in this study (a publicly traded parent company whose equity is priced by the market).

This problem has three manifestations. The first is subsidiary fragmentation: a single publicly traded parent may appear under dozens of distinct recipient names in USASpending. General Dynamics Corporation, for example, receives contracts through Electric Boat Corporation (submarines), Bath Iron Works (destroyers), NASSCO (auxiliary ships), General Dynamics Land Systems (armored vehicles), General Dynamics Information Technology (IT services), General Dynamics Mission Systems (electronics), and General Dynamics Ordnance and Tactical Systems (ammunition), among others. If these subsidiaries are treated as separate entities rather than consolidated into the parent, General Dynamics' apparent procurement exposure is dramatically understated, and several phantom 'firms' appear in the dataset.

The second manifestation is post-merger name persistence: when firms merge, the federal government's contracting records often continue to use the pre-merger legal names for years afterward because existing contract vehicles are not re-papered at the time of a merger. Raytheon Technologies (the 2020 merger of Raytheon Company and United Technologies Corporation) continued to receive contracts under RTX Corporation, Raytheon Technologies Corporation, Raytheon Intelligence & Space, Raytheon Missiles & Defense, and Rockwell Collins Inc. (a 2018 UTC acquisition) simultaneously. Without consolidation, these appear as five separate entities.

The third is legal name variation: minor differences in how the same entity's name is entered into the contracting system, punctuation, abbreviations, 'Inc.' vs 'Inc', comma placement, create apparent duplicates. BAE Systems Inc, BAE Systems INC, and BAE SYSTEMS INC. are the same entity but will appear as three separate rows in an API response.

### 3.4.2 Name Map Construction and Application

To address these problems, a name standardization procedure was applied in two passes. In the first pass, a Python dictionary (NAME_MAP) was constructed for each agency, mapping every known subsidiary name, legacy name, and typographic variant to a single canonical parent company name. This map was applied to the raw API response data immediately upon retrieval, before any aggregation or filtering. After the maps were applied, a groupby operation on the canonical name summed all subsidiary-level obligations into a single parent-level total. The maps were constructed by manually reviewing the top recipients for each agency and cross-referencing them against corporate ownership records, merger announcements, and the SEC's EDGAR database.

In total, approximately 150 unique subsidiary and variant names were mapped to their canonical publicly traded parents across the four agency-specific name maps. Representative examples are shown below.

| Raw USASpending Name(s) | Parent Company | Consolidation Type |
|---|---|---|
| Electric Boat Corp; Bath Iron Works; NASSCO; GD Land Systems; GD Information Technology; GD Mission Systems; GD Ordnance & Tactical Systems | General Dynamics Corporation | Wholly-owned operating subsidiaries |
| RTX Corporation; Raytheon Technologies Corporation; Raytheon Intelligence & Space; Raytheon Missiles & Defense; Rockwell Collins Inc. | Raytheon Company (RTX) | Post-merger entity, divisions, and acquired subsidiary |
| Northrop Grumman Systems Corporation; Northrop Grumman Systems Corp; Orbital Sciences LLC; ATK Launch Systems; Alliant Techsystems | Northrop Grumman Corporation | Operating subsidiaries and acquired companies (Orbital ATK 2018) |
| L3 Technologies Inc; Harris Corporation; L3Harris Technologies Inc.; L3Harris Integrated Systems L.P. | L3Harris Technologies | Post-merger entity (Jun 2019) and legacy names |
| Health Net Federal Services LLC | Centene Corporation | Centene acquired Health Net in 2016; subsidiary retained pre-acquisition name |

| | | |
|---|---|---|
| Express Scripts Inc | Cigna Corporation | Cigna acquired Express Scripts in 2018; subsidiary retained pre-acquisition name |
| Optumserve Technology Services; Optumserve Health Services | UnitedHealth Group Inc | Optum is UnitedHealth's health services subsidiary |
| Life Technologies Corporation; Fisher Scientific Company LLC; Thermo Electron North America | Thermo Fisher Scientific Inc | Acquired subsidiaries retained original federal contracting names |
| Vectrus Systems LLC; Vertex Aerospace LLC | V2X Inc | Vectrus and Vertex Aerospace merged in 2022 to form V2X |
| KBR Wyle Services LLC; KBR Services LLC | KBR Inc | Operating subsidiaries of KBR |
| CACI LLC; CACI Technologies LLC; CACI International Inc; CACI Federal Inc | CACI International Inc | Operating subsidiaries and legal name variants |

### 3.4.3 Second-Pass Consolidation and the Annual Pull Search Name Map

After the initial name mapping, a second consolidation pass was performed on the aggregated master firm list. The first-pass NAME_MAP was constructed by manually reviewing the top-ranked recipients for each agency, necessarily a finite review window. Subsidiary variants that appeared only in lower-ranked positions, and were therefore not encountered during that initial review, were not captured by the first-pass map. The second pass identified these residual cases by scanning the full aggregated list for names that were recognizably related to already-mapped parents but had not yet been consolidated. This included further CACI LLC variants, BAE Systems ship repair subsidiaries, and remaining General Dynamics ordnance subsidiaries.

A third name-matching challenge arose during the construction of the annual obligations panel, where contract obligations were retrieved separately for each firm x fiscal year x agency combination using a firm-name search against the USASpending API. This search is more sensitive to exact name matching than the original ranked-list pull. For 32 firms, the name used in the firm list (e.g., 'IBM CORPORATION') did not match the name under which the firm's contracts were recorded (e.g., 'INTERNATIONAL BUSINESS MACHINES CORPORATION'), resulting in zero obligations being returned for those firm-year combinations. A supplementary search name map (SEARCH_NAME_MAP) was constructed to provide the correct search term for each affected firm. Examples include: IBM Corporation → 'International Business Machines Corporation'; Cigna Corporation → 'Express Scripts Inc'; Humana Inc → 'Humana Government Business Inc'; Accenture PLC → 'Accenture Federal Services LLC'. The obligations pull was re-run for all 32 affected firms using the corrected search terms, and the results were merged back into the panel. After this correction, 95.7% of firm-year observations in the 2015-2024 regression window have non-zero exposure data.

The initial identification of subsidiary relationships, merger histories, and typographic variants was assisted by a large language model (Claude, Anthropic), with mappings subsequently verified by the authors against primary sources including SEC EDGAR filings and corporate merger announcements.

## 3.5 Firm Eligibility Filtering

### 3.5.1 Filtering Philosophy

The objective of the filtering stage was to retain only firms whose equity is publicly traded on a U.S. exchange or major OTC market. This requirement is non-negotiable because the study's dependent variables, stock returns, volatility, and beta, can only be computed for firms with observable market prices. The challenge is that the raw USASpending recipient lists contain a large and heterogeneous mix of entity types, the majority of which are not publicly traded. Of the top 1,000 DoD recipients, for example, fewer than 15% are publicly traded firms; the remainder are private companies, non-profit research institutions, universities, foreign firms, government entities, and joint ventures.

Two complementary filtering strategies were applied, assigned by the agency based on the structure of each agency's contractor base.

An important clarification on sequencing: the eligibility filter was applied at the level of the parent company, after the name standardization procedure described in Section 4 had already been applied. This ordering matters because many raw USASpending recipients are subsidiaries whose public-market status cannot be determined by looking up the subsidiary name alone, for example 'Electric Boat Corporation' does not appear on any stock exchange, but its parent General Dynamics does. The question 'is this recipient publicly traded?' can therefore only be answered reliably after subsidiary names have been resolved to their parents. Once that resolution was complete, the eligibility check was more straightforward: each parent was verified against stock exchange listings, SEC EDGAR filings, and CRSP coverage. The final sample comprises 118 publicly traded parent companies.

### 3.5.2 Inclusion-Based Filtering: HHS, NASA, and VA

For HHS, NASA, and the VA, only firms that had been explicitly pre-verified as publicly traded were retained. Every other recipient was dropped, regardless of contract size. This was necessary because these agencies award contracts to contractor bases that are overwhelmingly composed of non-public entities, hospitals, universities, non-profits, foreign suppliers, small businesses, and state agencies vastly outnumber the publicly traded firms among their top recipients. It would not be feasible to enumerate all the non-public entities to block; it is far more tractable to enumerate the public firms to keep.

The list of publicly traded firms was constructed by manually reviewing the top 200-300 recipients for each agency, cross-referencing each against stock exchange listings, SEC filings, and Bloomberg, and recording those with active U.S. exchange listings or significant ADR programs. The allowlists comprised approximately 55 firms for HHS, 40 for NASA, and 50 for the VA. These lists were then used as the sole filter, any recipient not on the allowlist was automatically excluded.

### 3.5.3 Exclusion-Based Filtering: DoD

On the contrary, for the Department of Defense, all recipients were initially retained and then a curated list of confirmed non-public entities was explicitly removed. This approach was viable for DoD because the agency's contractor base is substantially more concentrated among publicly traded defense primes than the other agencies; the top 50 DoD contractors by dollar value are predominantly large, well-known public companies.

The DoD exclusion list included entities in the following categories:

- Federally Funded Research and Development Centers (FFRDCs): these are non-profit entities sponsored by federal agencies to perform research that the government cannot procure commercially. Examples include MIT Lincoln Laboratory, Johns Hopkins APL, The Aerospace Corporation, MITRE Corporation, RAND Corporation, Draper Laboratory, the Institute for Defense Analyses, and CNA. Although they receive large DoD contracts, they have no publicly traded equity.
- Non-profit research universities: Penn State Applied Research Laboratory, Carnegie Mellon, Georgia Tech Research Institute, UT Austin, USC, University of Washington, and University of Dayton all appear in DoD contract data for sponsored research but are non-profit educational institutions.
- Private defense contractors: SpaceX, Sierra Nevada Corporation, General Atomics, Bechtel Corporation, Peraton, and AM General are significant DoD contractors but are privately held.
- Alaska Native Corporations (ANCs) and AbilityOne contractors: ASRC Federal, Akima, Alutiiq, Chugach Federal Solutions, Cherokee Nation entities, and Bowhead receive DoD contracts under set-aside programs but are not publicly traded.
- Joint ventures with no standalone equity: entities such as Longbow LLC (Boeing/Raytheon JV for Apache helicopters), CFM International (GE/Safran JV for jet engines), and Data Link Solutions (Collins/BAE JV) are legal entities that receive contracts but have no independently traded equity.

- Foreign-listed firms with no U.S. listing: Rolls-Royce (LSE), Austal (ASX), Airbus (Euronext Paris), Kongsberg (Oslo), Thales (Euronext Paris), Saab (Stockholm), and Serco (LSE) receive DoD contracts but are not listed on U.S. exchanges and are not covered by CRSP.
- Firms that went private during the sample period: ManTech International was taken private by The Carlyle Group in July 2022, Atlas Air was acquired in August 2023, and Maxar Technologies was acquired by Advent International in January 2023. These firms have partial CRSP coverage and are retained in the panel for the years they were public but are flagged in the notes column.

### 3.5.4 Final Ticker Filter and Quality Control

As a final quality control step, each remaining firm was assigned its primary exchange ticker symbol and the master list was filtered to retain only rows where a ticker had been successfully assigned. This ensured that any entity that had slipped through the inclusion or exclusion filters but lacked a verifiable public market listing was automatically excluded. To clarify the overall sequence: name standardization was applied first, consolidating all subsidiary-level recipients into canonical parent companies. Eligibility filtering was then applied to those canonical parents, not to individual subsidiaries. The ticker assignment described here is therefore the final confirmation step in that pipeline, not the primary filter. The ticker also served as the key for the subsequent WRDS lookup, tickers were used to retrieve GVKEYs from Compustat, which were then used to retrieve PERMNOs from the CCM link table.

One correction was required at this stage: AmerisourceBergen Corporation rebranded to Cencora in 2023 and changed its ticker from ABC to COR. Compustat stores the firm under the new name; the ticker in the firm list was updated to COR before the WRDS lookup to ensure the match succeeded. The final ticker filter reduced the master list from approximately 650 consolidated entries to 120 confirmed publicly traded firms, of which 118 received PERMNOs (BAE Systems and Roche Holding had no CRSP coverage and were excluded from the analysis panel).

## 3.6 Panel Construction

### 3.6.1 Identifier Lookup

GVKEYs were retrieved programmatically from the Compustat name history table (comp.names) by matching on ticker symbol. GVKEYs are stored as 6-character zero-padded strings in Compustat (e.g., ‘006774’ for Lockheed Martin), but Python reads them as integers from CSV files, dropping the leading zeros. This caused the WRDS SQL query to find no matches for the affected firms. The fix was to apply Python’s str.zfill(6) function to pad all

GVKEYs back to six digits before passing them to the query. After this correction, 118 of 120 firms were matched to GVKEYs, and all 118 were subsequently matched to PERMNOs via the CCM link table.

### 3.6.2 Timing Alignment and Lag Structure

The three datasets cover different time periods and different time frequencies. Aligning them correctly is critical for avoiding look-ahead bias, the error of using information that was not yet available at the time a return was earned. The alignment used in this study follows the standard approach in the accounting and finance literature and proceeds as follows.

USASpending records contract obligations on a U.S. federal fiscal year basis (October 1 - September 30). Compustat records financials on each firm's own fiscal year, which for most firms in the sample ends in December but varies. CRSP records returns by calendar month. The panel merges these on a (ticker, year) key using the following lag structure:

| Variable | Source Period | Matched To |
| --- | --- | --- |
| Contract obligations | Federal FY t-1 (Oct t-2 - Sep t-1) | Compustat FY t revenue (denominator) |
| Revenue and controls | Compustat FY t | CRSP calendar year t+1 monthly returns |
| Monthly returns | CRSP calendar year t+1 | Exposure ratio constructed from FY t-1 obligations / FY t revenue |

Put simply: to explain what a firm's stock does in 2020, we use contract obligations from 2018 and revenue from 2019, as this is information that any investor could have looked up before January 2020 began. The years shift by one each time the return year changes: for 2021 returns, we use 2019 obligations and 2020 revenue; for 2019 returns, we use 2017 obligations and 2018 revenue. The formal notation in the table above captures this same relationship, t-1 obligations and t revenue explaining t+1 returns, but the underlying logic is simply that everything used to construct the exposure measure is already in the past by the time the returns being explained are earned.

### 3.6.3 Exposure Variable Construction

The primary independent variable, procurement exposure, is defined as:

**Exposure(i,t) = Obligations(i, FY t-2) / Revenue(i, FY t-1)**

In this notation, t refers to the return year. For a firm with a December fiscal year end and a return year of 2020 (t=2020), obligations come from federal fiscal year 2018 (t-2) and revenue from Compustat fiscal year 2019 (t-1). To make this concrete, the following traces that single observation through all three sources:

- Contract obligations: federal fiscal year 2018 (October 1, 2017 - September 30, 2018)
- Revenue and controls: Compustat fiscal year 2019 (January 1 - December 31, 2019)
- Monthly returns: CRSP calendar year 2020 (January - December 2020)

The two-year gap between the obligations period and the return year is not a modelling choice but a consequence of the three-source timing structure: federal fiscal years end in September, firm fiscal years end in December, and returns run on calendar years. At no point does the exposure measure use information from 2020 itself, both the numerator and denominator were publicly observable before the first return in the window was earned.

This ratio measures the fraction of firm revenue attributable to government contracts in the prior fiscal year. A ratio of 0.50 means the firm received contract obligations equal to half its annual revenue in the prior year. The ratio uses revenue (net sales) rather than total assets as the denominator because it directly measures the revenue-generating relationship between government contracts and firm operations, whereas assets are a stock variable that would dilute the measure for capital-intensive firms.

The raw exposure ratio is highly right-skewed: most firms in the sample have ratios below 0.10, reflecting the fact that government contracts represent one of several revenue streams for the majority of firms in the sample. Large diversified firms such as Microsoft, Amazon, and Dell Technologies appear in the sample through substantial DoD and HHS information technology contracts, but generate the overwhelming majority of their revenues from commercial cloud services, enterprise software, and hardware sales. Similarly, large pharmaceutical firms such as Pfizer and Johnson & Johnson receive significant HHS procurement obligations, particularly during the COVID-19 period, but derive most of their revenue from commercial drug sales across global markets. For these firms, government procurement is economically meaningful but represents a small share of total revenue, producing exposure ratios well below 0.10. At the other end of the distribution, a small number of pure-play contractors, like Huntington Ingalls, GLDD, V2X, and SIGA Technologies, who derive the substantial majority of their revenues directly from federal contracts, producing ratios approaching or exceeding 1.0 in some years. To prevent these outliers from dominating regression estimates, the ratio was winsorized at the 1st and 99th percentile. A log transformation, log(1 + exposure_ratio), was then applied to further compress the distribution and improve the linearity assumption for OLS. The distribution of the winsorized exposure ratio in the final panel has a mean of 0.095, a median of 0.005, and a standard deviation of 0.213.

# 4. RESULTS

## 4.1 Descriptive Statistics

Table 1 presents summary statistics for the full regression sample, which comprises 14,104 firm-month observations across 118 publicly traded U.S. firms over the period 2014 to 2024. The average monthly stock return in the sample is 1.10%, with a standard deviation of 8.80%, reflecting the predominantly bull market conditions of the sample period. Following standard practice in the empirical asset pricing literature, monthly returns are winsorized at the 1st and 99th percentile, in order to bound the distribution between -23.7% and +29.6%, which prevents extreme observations from distorting regression estimates. The winsorization bounds reflect the empirical return distribution of the sample, which is driven largely by firm-specific binary events among smaller biotech and defence-adjacent firms, for example, Moderna's +126% return in November 2020 following its covid vaccine announcement and Chimerix's -78% collapse in December 2015 after a failed drug trial, events that, while economically real, would disproportionately distort regression estimates if left unaddressed.

The central independent variable, log exposure, has a mean of 0.074 and a median of 0.003, confirming that this is a right-skewed distribution. The median winsorized exposure ratio of 0.004 indicates that for the typical firm in the sample, government contract obligations represent less than one percent of annual revenue. This reflects the fact that the sample includes many large and diversified firms, such as Microsoft, Amazon, and Pfizer, for whom federal procurement is economically meaningful but constitutes a small share of total revenues. At the other end of the distribution, pure-play contractors such as Huntington Ingalls, V2X, and Booz Allen Hamilton have exposure ratios approaching or exceeding 1.0 in some years.

The control variables behave as expected. The mean leverage ratio is 0.25, consistent with moderate financial risk across the sample. The R&D ratio has a mean of 2.85 but a standard deviation of 84.2 and a maximum of 2,877, reflecting a severely right-skewed distribution driven by a small number of biotech firms where R&D expenditure temporarily exceeds total revenue. While the coefficient on rd_ratio in the regressions is sufficiently small that this does not materially affect the results, this distributional feature is acknowledged as a limitation and is addressed in the robustness discussion. The agency flags indicate that 66% of observations belong to DoD-contracting firms, 48% to HHS, 26% to NASA, and 16% to VA, with firms able to appear in multiple agency categories simultaneously.

## 4.2 Main Return Regressions

Table 2 presents the main panel regression results, structured as three progressively richer specifications, M1, M2, and M3. These are all estimated using two-way fixed effects with firm and year fixed effects and standard errors clustered by firm. The dependent variable in all specifications is the monthly winsorized stock return. The progressive model design serves an important methodological purpose: if the coefficient on log exposure remains stable as controls are added, this provides evidence that the relationship is not driven by omitted correlated variables. If instead the coefficient moves, it signals that the raw correlation in the simpler model was driven by confounding firm characteristics rather than procurement exposure itself.

M1 is the baseline specification, regressing monthly returns on log exposure alone, with firm and year fixed effects. The firm fixed effects absorb all time-invariant firm characteristics, industry membership, business model, geographic concentration, that might be correlated with both procurement exposure and returns. The year fixed effects absorb all aggregate time-series variation, including broad market movements, interest rate changes, and macroeconomic shocks that affect all firms simultaneously. The coefficient on log exposure in M1 is +0.0263, with a standard error of 0.0159. This positive but statistically insignificant coefficient suggests a weak raw association between procurement exposure and returns that does not meet conventional significance thresholds.

M2 adds log total assets and leverage as controls. These two variables are the most important potential confounders in the specification. Larger firms may systematically differ in their procurement relationships and in their return profiles, while more highly leveraged firms carry greater financial risk that is independently priced by the market. Including these controls is standard in empirical asset pricing studies of firm-level returns. The coefficient on log exposure falls to +0.0172 (SE = 0.0141), suggesting that part of the positive raw correlation in M1 reflected the fact that high-exposure firms tend to be larger and more leveraged. Both controls are highly significant, with log assets negative (-0.0166, $p < 0.01$) and leverage positive (+0.0293, $p < 0.01$), consistent with the standard empirical asset pricing literature.

M3 is the preferred full specification, adding log market cap and the R&D ratio on top of M2. Log market cap is included as a second size measure that captures market-based valuation rather than accounting-based assets, and is updated monthly from CRSP rather than annually from Compustat, making it a more timely control variable. The R&D ratio controls for growth optionality and innovation intensity, which may independently affect both a firm's likelihood of winning government contracts and its return profile. In M3, the coefficient on log exposure is -0.0047 (SE = 0.0137), effectively zero. The full set of controls absorbs all systematic variation in returns that was previously attributed to procurement exposure in M1 and M2, leaving no residual predictive power for the exposure variable in the full sample.

This null result in the full sample is itself informative. It suggests that on average, across all four agencies and all firm types, the market prices government procurement exposure efficiently, and there is no systematic under- or over-pricing of procurement-dependent firms in the cross-section.The controls behave as expected throughout, log market cap is positive and strongly significant (+0.0343, $p < 0.01$), leverage is positive and significant (+0.0637, $p < 0.01$), log assets is negative and significant (-0.0329, $p < 0.001$). The within $R^2$ increases from 0.06% in M1 to 2.47% in M3, reflecting the expanded explanatory power added by the controls. The relatively modest $R^2$ is expected in monthly return regressions with fixed effects, where the majority of cross-sectional variation is absorbed by the firm and year dummies.

## 4.3 Agency Subsample Regressions

Table 3 presents the agency subsample results, running the full M3 specification separately for firms contracting with each of the four agencies. This is the central empirical contribution of the thesis. The existing literature, Goldman (2020), Parajuli (2023), Dhaliwal et al. (2016), treat federal procurement as a homogeneous category. The agency subsample design tests whether this aggregation conceals meaningful heterogeneity in how the market prices different types of government revenue. The four agencies differ fundamentally in their contractual structures, budget stability, and sensitivity to political and macroeconomic shocks, as documented in the literature review, and these differences should in theory produce different return implications if the market prices them correctly.

For DoD-contracting firms (76 firms, 9,299 observations), the coefficient on log exposure is -0.0087 (SE = 0.0170). This null result indicates that the market prices DoD procurement exposure efficiently on average; there is no systematic return premium or discount associated with higher DoD dependence. This finding is somewhat surprising given our theoretical expectation that DoD's long-term multiyear procurement mechanisms should reduce cash flow uncertainty and thus lower required returns. One interpretation, consistent with the institutional evidence reviewed in the literature, is that DoD acquisition programs are sufficiently subject to cost overruns, schedule delays, and technological uncertainty, as documented repeatedly by the Government Accountability Office. It would then follow that the revenue stability benefit of long-term contracts is offset by execution risk, leaving no net effect on required returns.

For HHS-contracting firms (54 firms, 6,818 observations), the coefficient is +0.0014 (SE = 0.0179). This is essentially zero, indicating no relationship between HHS procurement exposure and stock returns. HHS procurement spans a heterogeneous set of firm types, including pharmaceutical manufacturers, managed care organizations, health IT firms, and laboratory services providers, which may partially explain the absence of a systematic effect, as the return implications of HHS contracts likely vary substantially across these very different business models.

For NASA-contracting firms (34 firms, 3,718 observations), the coefficient on log exposure is -0.0511 (SE = 0.0295, $p < 0.1$), which is marginally significant at the 10% level. Firms with greater exposure to NASA procurement earn meaningfully lower monthly returns, with a one-unit increase in log exposure associated with a 5.1 percentage point reduction in monthly returns. NASA's contracting structure, which is characterized by long-term cost-plus agreements concentrated among a small number of specialized aerospace firms, provides a particularly stable and predictable revenue stream that the market correctly prices as lower risk, and therefore demands lower returns. The result is further supported by the behavior of the R&D ratio control in the NASA subsample, which is negative and highly significant (-0.1739, $p < 0.01$), suggesting that R&D-intensive NASA contractors earn particularly low returns, consistent with the interpretation that cost-plus contracts effectively transfer innovation risk from the contractor to the government, reducing the firm-specific risk premium that R&D intensity would otherwise command.

For VA-contracting firms (19 firms, 2,311 observations), the coefficient is -0.0219 (SE = 0.0366). The direction is consistent with NASA, negative, suggesting a return discount for VA-exposed firms, but the result is not statistically significant. This should be interpreted with caution given the small number of firms in the VA subsample, which substantially limits statistical power. The VA result cannot be used to draw conclusions about whether the market prices VA procurement differently from other agencies.

Taken together, the agency subsample results support the thesis's central argument that aggregate procurement exposure masks important cross-agency heterogeneity. The market does not price all government revenue equally, NASA procurement exposure is associated with a marginally significant return discount consistent with efficient pricing of superior cash flow stability, while DoD, HHS, and VA show no significant effects.

## 4.4 Interaction Regressions

Table 4 presents two interaction specifications, each motivated by established findings in the government procurement and asset pricing literature. Both use the full M3 control set and two-way fixed effects.

The first specification, reported in Table 4a, adds an interaction between log exposure and a COVID-period dummy equal to one for the months of March 2020 through December 2021. This interaction is motivated by the structural shift in federal procurement during the pandemic period, particularly the acceleration of HHS spending on vaccine and therapeutic contracts, and tests whether the return implications of procurement exposure changed during this period. The main coefficient on log exposure is +0.0034, consistent with the M3 null result. The interaction coefficient is -0.0196 (SE = 0.0110), marginally significant at the 10% level. This negative coefficient indicates that during the COVID period, higher procurement exposure was associated with lower returns relative to the non-COVID baseline, the opposite of what might be expected if government contracts provided a safe haven during the economic disruption of 2020. A plausible interpretation is that the broader equity market during 2020-2021 was dominated by a risk-on rotation into growth and technology stocks, which rewarded uncertainty and optionality rather than cash flow stability. In this environment, the predictable revenue streams associated with government contracts became a relative disadvantage from a return perspective, as investors assigned lower multiples to stable but slower-growing procurement-dependent firms. This finding, while only marginally significant, is consistent with the broader literature on how market sentiment affects the pricing of cash flow stability.

The second specification, reported in Table 4b, adds an interaction between log exposure and a Democratic-presidency dummy, following the methodology of Belo et al. (2013), who document that government-exposed firms earn significantly higher returns during Democratic administrations. In the present sample, the Democrat dummy equals one during the Obama administration through January 2017 and the Biden administration from January 2021 onwards, and zero during the first Trump administration. The main coefficient on log exposure is +0.0031. The interaction coefficient is -0.0180 (SE = 0.0119), negative but not statistically significant. The direction is noteworthy, as it is the opposite of Belo et al. 's (2013) finding, suggesting that in this sample, high-exposure firms earn slightly lower returns during Democratic presidencies rather than higher. However, with only 11 years of data spanning parts of three administrations, the sample provides insufficient statistical power to draw conclusions about political cycle effects. The result is reported for completeness and as a point of comparison with the existing literature, but cannot be interpreted as evidence against Belo et al.'s (2013) findings given the sample size constraints.

## 4.5 Fama-MacBeth Robustness

As a robustness check on the main panel findings, Table 5 presents Fama-MacBeth cross-sectional regression results following Fama and MacBeth (1973). The procedure runs a separate cross-sectional OLS regression of compounded annual returns on log exposure and controls for each year in the sample, averaging the resulting coefficients across 11 annual cross-sections. This addresses potential concerns about serial correlation in the pooled panel estimates. The average coefficient on log exposure is +0.0818 (SE = 0.0697, t-stat = 1.17), positive but not statistically significant. The direction is consistent with the M1 panel result and the larger magnitude reflects the use of annual rather than monthly returns. The absence of significance across only 11 cross-sections is expected, as Fama-MacBeth standard errors increase as the number of time periods decreases, limiting power to detect effects of the magnitude documented in the panel. Crucially, the Fama-MacBeth results do not contradict the main findings and the NASA subsample result remains the primary conclusion of the analysis.

# 5. LIMITATIONS

Several limitations of the study should be acknowledged when interpreting the results.

The most fundamental is the obligation timing mismatch between USASpending and Compustat. USASpending records contract obligations at the point of award, while Compustat records revenue at the point of recognition. For long-duration platform contracts, a shipbuilding contract awarded in 2018 may generate Compustat revenue through 2026, there is a systematic lag between the obligation year and the fiscal year in which the associated revenue appears. This creates measurement error in the exposure ratio that will attenuate regression coefficients toward zero, meaning the true relationship between procurement exposure and returns may be stronger than the estimates reported here suggest.

A second limitation concerns the top-recipient sampling design. By constructing the firm list from the top USASpending recipients by agency, the sample overrepresents large, financially stable contractors and is not representative of the full universe of firms with any government revenue. Results should not be extrapolated beyond the population of significant procurement counterparties studied here.

Third, residual name matching noise affects a small number of firms despite the three-pass standardization procedure. Four managed care firms, Cigna, Centene, Humana, and UnitedHealth, show near-zero exposure in parts of the panel, likely reflecting the genuine difficulty of matching health insurance premium payments to firm-name searches in USASpending. These firms' exposure ratios may understate their true government revenue dependence, introducing attenuation bias in both the HHS and VA subsample results.

Fourth, the COVID period created non-stationary exposure ratios for several HHS firms. Moderna, Pfizer, and Emergent BioSolutions experienced exposure ratio spikes during 2020-2022 driven by emergency procurement that is unlikely to reflect stable long-run procurement relationships. While the COVID interaction specification in Table 4a partially addresses this by allowing the exposure effect to differ during the pandemic period, the HHS subsample results should be interpreted with particular caution.

Finally, the unbalanced panel includes nine firms with fewer than 60 months of CRSP coverage due to late IPOs or early delistings, primarily SPAC listings from 2020-2023. While their inclusion does not materially affect the full-sample results, the agency subsample estimates, particularly VA with only 19 firms, are sensitive to partial-coverage firms and should be interpreted accordingly.

# 6. FINAL CONCLUSIONS

This study set out to examine whether a firm's dependence on federal contract revenue is reflected in its equity returns, and whether that relationship varies meaningfully across different government agencies. The full-sample results offer a clear answer to the first question: on average, procurement exposure carries no systematic return premium or discount. Once firm size, leverage, and market capitalization are controlled for, the coefficient on log exposure is effectively zero across the 118-firm, 14,104-observation panel. This null result is itself informative. It is consistent with the efficient markets interpretation of the literature, in other words, if government contracts reduce earnings uncertainty, that reduction should be priced into valuations rather than generating persistent abnormal returns. The market, in aggregate, appears to do exactly that.

The agency subsample results provide a richer and more nuanced picture. Disaggregating by contracting agency reveals heterogeneity that the full-sample specification conceals. The most notable finding is for NASA-contracting firms, where a one-unit increase in log exposure is associated with a 5.1 percentage point reduction in monthly returns, significant at the 10% level. This result is consistent with the interpretation that NASA's long-term, cost-plus contracting structure provides an unusually stable and predictable revenue stream, one the market correctly prices as lower risk and therefore requiring lower returns. The behavior of the R&D ratio in the NASA subsample reinforces this reading: the negative and highly significant coefficient suggests that even R&D-intensive NASA contractors earn lower returns, consistent with cost-plus contracts transferring innovation risk from contractor to government. By contrast, DoD, HHS, and VA subsamples show no significant effects, though the directions vary, and the VA result in particular is difficult to interpret given the small number of firms and limited statistical power.

Taken together, these findings advance procurement literature in two ways. First, they provide some evidence against the common assumption that federal contract revenue is a homogeneous exposure category. The market appears to price NASA procurement differently from DoD or HHS procurement, suggesting that contract structure and revenue predictability, rather than merely having the government as counterparty, are what matter for risk pricing. Second, the interaction results offer tentative support for the idea that the return implications of procurement exposure are time-varying: the marginally significant negative COVID interaction suggests that the revenue stability associated with government contracts became a relative disadvantage during the 2020-2021 risk-on market rotation, when investors rewarded growth and optionality over predictability. Both of these findings point toward directions for future research with longer time series, broader firm coverage, and contract-level data that would allow more precise measurement of the mechanism.

# 7. REFERENCES


Belo, Frederico; Gala, Vito D; Li, Jun, 2013, Government spending, political cycles, and the cross section of stock returns, Journal of Financial Economics, Vol.107 (2), p.305-324

Centers for Medicare & Medicaid Services, 2026, Contracting with CMS, https://www.cms.gov/about-cms/work-with-us/contracting

Cortes, Gustavo S; Vossmeyer, Angela; Weidenmier, Marc D, 2022, Stock volatility and the war puzzle: the military demand channel, NBER Working Paper 29837

Dhaliwal, Dan; Judd, J. Scott; Serfling, Matthew; Shaikh, Sarah, 2016, Customer concentration risk and the cost of equity capital, Journal of Accounting & Economics, Vol. 61 (1), p.23-48

Diltz, J. David, 1990, Valuation effects of government contract awards, Quarterly Journal of Business and Economics, Vol. 29 (3), p.3-15

Eibner, Christine; Krull, Heather; Brown, Kristine M.; Cefalu, Matthew; Mulcahy, Andrew W.; Pollard, Michael; Shetty, Kanaka; Adamson, David M.; Amaral, Ernesto F. L.; Armour, Philip; Beleche, Trinidad; Bogdan, Olena; Hastings, Jaime; Kapinos, Kandice; Kress, Amii; Mendelsohn, Joshua; Ross, Rachel; Rutter, Carolyn M.; Weinick, Robin M.; Woods, Dulani; Hosek, Susan D.; Farmer, Carrie M. 2016, Current and projected characteristics and unique health care needs of the patient population served by the Department of Veterans Affairs, RAND HealthQ 5(4):13, https://pmc.ncbi.nlm.nih.gov/articles/PMC5158228/

Fama, Eugene F, 1970, Efficient capital markets: a review of theory and empirical work, The Journal of Finance (New York), Vol. 25 (2), p.383 - 417

Goldman, Jim, 2020, Government as customer of last resort: The Stabilizing Effects of Government Purchases on Firms, The Review of Financial Studies, Vol. 33 (2), p.610-643

Government Accountability Office, 2025a, Weapon systems annual assessment: DoD leaders should ensure that new programs are structured for speed and innovation, GAO-25-107569, https://www.gao.gov/assets/gao-25-107569.pdf

Government Accountability Office, 2025b, NASA: Assessments of major projects, GAO-25-107591, https://www.gao.gov/assets/gao-25-107591.pdf

Larson, Stephen J and Picou, Armand, 2007, The market's response to contract award announcement: government versus corporate contracts, Journal of Public Procurement, Vol. 2 (2), p.237-251

Lindbergh, Rachel, 2026, NASA Appropriations and authorizations: At a glance, R43419, https://www.congress.gov/crs-product/R43419

Neenan, Alexndra G., 2026, Defense primer: Procurement, IF10599, https://www.congress.gov/crs-product/IF10599

O'Rourke, Donald, 2025, Multiyear procurement (MYP) and block buy contracting in defense acquisition: Background and issues for congress, R41909, https://www.congress.gov/crs-product/R41909

Panangala, Sidath Viranga; Sussman, Jared S., 2024, Department of Veterans Affairs FY2024 appropriations, R48056, https://www.congress.gov/crs-product/R48056

Parajuli, Bharat Raj, 2023, Asset pricing implications of firm's government sales dependency, Review of Asset Pricing Studies, Vol.13 (1), p.146-180

Pyun, Chaehyun, 2025, Trading on government contracts: the investment potential of public procurement contracts, Economic Letters, Vol. 252, article.112335, p.112335

Samunderu, Eyden and Anita Yordanova, 2024, The effect of public contract winning announcements on share prices: an event-based study on the pharmaceutical industry, Journal of Economic Analysis, Vol. 3 (1), p.1-26

Sekar, Kavya; Sheikh, Hassan Z., 2025, Pandemic and all-hazards preparedness act: An overview, IF12683, https://www.congress.gov/crs-product/IF12683

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, The Journal of Finance (New York), Vol.19 (3), p.425

# 8. APPENDIX

## 8.1 Appendix A - Variable Construction

**A.1 Dependent Variable**

Monthly stock return, compounded to include dividends and adjusted for splits, retrieved from the CRSP Monthly Stock File. Where a delisting return exists in the same month, returns are compounded as:

- r(i,t) = (1 + ret(i,t)) × (1 + dlret(i,t)) - 1

Returns are winsorized at the 1st and 99th percentile of the cross-sectional distribution within the regression sample.

**A.2 Primary Independent Variable**

The procurement exposure ratio is defined as:

- Exposure(i,t) = Obligations(i, FY t-2) / Revenue(i, FY t-1)

where Obligations are total federal contract obligations from USASpending in federal fiscal year t-2 (ending September 30), Revenue is net sales from Compustat in firm fiscal year t-1, and t refers to the CRSP return year. The raw ratio is winsorized at the 1st and 99th percentile and log-transformed:

- log_exposure(i,t) = log(1 + Exposure(i,t)_w)

where the subscript w denotes the winsorized value. Negative obligation values (contract de-obligations) are floored at zero before transformation.

**A.3 Control Variables**

**Log total assets:**

- log_assets(i,t) = log(AT(i,t))

where AT is total assets from Compustat (item AT), measured in millions of USD.

**Leverage:**

- leverage(i,t) = LT(i,t) / AT(i,t)

where LT is total long-term debt and AT is total assets, both from Compustat.

**R&D ratio:**

- rd_ratio(i,t) = XRD(i,t) / AT(i,t)

where XRD is research and development expenditure from Compustat. Firms reporting zero or missing R&D are assigned a value of zero, consistent with the convention that non-reporting firms have zero R&D spend.

**Log market capitalisation:**

- log_mktcap(i,t) = log(PRC(i,t) × SHROUT(i,t))

where PRC is the month-end closing price and SHROUT is shares outstanding, both from the CRSP Monthly Stock File. Market capitalisation is updated monthly.

**A.4 Binary Indicators**

**COVID period:**

- covid(t) = 1 if March 2020 ≤ t ≤ December 2021, else 0

**Democrat presidency:**

- democrat(t) = 1 if t < January 20 2017 or t ≥ January 20 2021, else 0

corresponding to the Obama administration through January 2017 and the Biden administration from January 2021, with the Trump administration (January 2017 - January 2021) as the reference period.

**Agency flags:**

is_dod(i) = 1 if firm i received positive obligations from the Department of Defense in at least one fiscal year in the sample, else 0

Equivalent definitions apply for is_hhs, is_nasa, and is_va.

**A.5 Coverage Statistics**

CRSP return coverage is defined as the share of regression sample observations with a non-missing monthly return after delisting return merging and winsorization:

- CRSP coverage = |{(i,t) : ret(i,t) is non-missing}| / |regression sample| = 14,104 / 14,104 = 100%

Exposure ratio coverage is defined as the share of regression sample observations with a non-missing winsorized exposure ratio:

- Exposure coverage = |{(i,t) : exposure_ratio_w(i,t) is non-missing}| / |regression sample| = 13,564 / 14,104 = 96.2%

The 3.8% of observations with missing exposure ratios arise primarily from firm-years where Compustat revenue is missing or zero, making the denominator of the exposure ratio undefined.

## 8.2 Appendix B - Script Architecture & Results

### 8.2.1 Script Architecture

The full data pipeline consists of seven Python scripts, each taking the output of the previous as input. All scripts are archived in the project repository.

| Script | Input | Output | Key Operations |
|---|---|---|---|
| `step1_fetch_agencies.py` | `USASpending API` | `Master_firm_list_v2.csv, *_checkpoint.csv` | API pull with retry/checkpoint; first-pass NAME_MAP application; inclusion/exclusion filtering |
| `step2_filter_firms.py` | `master_firm_list_v2.csv` | `master_firm_list_v3.csv` | Second-pass name consolidation; ticker pre-fill; notes and flags |
| `step3_match_identifiers.py` | `master_firm_list_THESIS.csv (manual), WRDS API` | `master_firm_list_GVKEYS.csv` | Final subsidiary consolidation; ticker filter; DROP/NOISE/ADR flags |
| `step4_pull_data.py` | `master_firm_list_GVKEYS.csv,WRDS API` | `crsp_monthly.csv, compustat_annual.csv` | GVKEY lookup via comp.names; PERMNO lookup via CCM; GVKEY zero-padding fix |
| `step5_build_panel.py (reconstructed)` | `crsp_monthly.csv, compustat_annual.csv, master_firm_list_GVKEYS.csv, USASpending API` | `analysis_panel.csv,usaspending_obligations.csv` | CRSP monthly returns + delisting; Compustat North America + Global fundamentals |
| `step6_fix_obligations.py` | `Analysis_panel.csv, usaspending_obligations.csv, USASpending API` | `analysis_panel_fixed.csv, usaspending_obligations_fixed.csv` | USASpending annual pull; SEARCH_NAME_MAP correction; exposure construction; panel merge |
| `step7_regressions.py` | `analysis_panel_fixed.csv` | `regression_results.txt, table1_summary.tex,` | Return winsorization; two-way FE panel regressions (M1/M2/M3); agency |

| | | | |
|---|---|---|---|
| | | `table2_main.tex, table3_agency.tex, table4_interactions.tex, table5_fmb.tex` | subsamples (DoD/HHS/NASA/VA); COVID and political cycle interactions; Fama-MacBeth cross-sectional robustness |

### 8.2.2 Results

**Table 1: Summary Statistics**

| Variable | N | Mean | Std Dev | P25 | Median | P75 | Min | Max |
|---|---|---|---|---|---|---|---|---|
| | | | | *Returns and Exposure* | | | | |
| Monthly return | 14,104 | 0.011 | 0.088 | -0.039 | 0.010 | 0.058 | -0.237 | 0.296 |
| Log exposure | 14,104 | 0.074 | 0.153 | 0.000 | 0.003 | 0.049 | 0.000 | 0.780 |
| Exposure ratio (winsorized) | 13,564 | 0.095 | 0.212 | 0.000 | 0.004 | 0.057 | 0.000 | 1.181 |
| | | | | *Firm Characteristics* | | | | |
| Log market cap | 14,104 | 9.701 | 1.918 | 8.386 | 9.992 | 11.128 | 3.641 | 15.016 |
| Log total assets | 14,104 | 9.640 | 1.933 | 8.195 | 9.934 | 11.201 | 0.251 | 13.221 |
| Leverage | 14,104 | 0.252 | 0.149 | 0.152 | 0.248 | 0.338 | 0.000 | 2.261 |
| R&D ratio | 14,104 | 2.852 | 84.165 | 0.000 | 0.024 | 0.076 | 0.000 | 2877.9 |
| | | | | *Indicator Variables* | | | | |
| COVID period | 14,104 | 0.279 | 0.448 | 0.000 | 0.000 | 1.000 | 0 | 1 |
| Democrat administration | 14,104 | 0.641 | 0.480 | 0.000 | 1.000 | 1.000 | 0 | 1 |
| DoD firm | 14,104 | 0.659 | 0.474 | 0.000 | 1.000 | 1.000 | 0 | 1 |
| HHS firm | 14,104 | 0.483 | 0.500 | 0.000 | 0.000 | 1.000 | 0 | 1 |
| NASA firm | 14,104 | 0.264 | 0.441 | 0.000 | 0.000 | 1.000 | 0 | 1 |
| VA firm | 14,104 | 0.164 | 0.370 | 0.000 | 0.000 | 0.000 | 0 | 1 |

Notes: Sample covers 118 firms, 2014-2024. Exposure ratio winsorized at 1st and 99th percentiles. Log exposure = log(1 + exposure ratio)

**Table 2: Procurement Exposure and Monthly Stock Returns**

| Variable | M1 | M2 | M3 |
|---|---|---|---|
| | Baseline | + Controls | Full Model |
| Log exposure | 0.0263 | 0.0172 | -0.0047 |
| | (0.0159) | (0.0141) | (0.0137) |
| Log market cap | | | 0.0343*** |
| | | | (0.0026) |
| Leverage | | 0.0293*** | 0.0637*** |
| | | (0.0106) | (0.0104) |
| R&D ratio | | | 0.0000** |
| | | | (0.0000) |
| Log total assets | | -0.0166*** | -0.0329*** |
| | | (0.0024) | (0.0036) |
| Firm FE | Yes | Yes | Yes |
| Time FE | Yes | Yes | Yes |
| Observations | 14,104 | 14,104 | 14,104 |
| $R^2$ (within) | 0.0006 | 0.0039 | 0.0247 |

Notes: Dependent variable is monthly stock return. Two-way fixed effects panel regression with firm and time fixed effects. Standard errors clustered by firm in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

**Table 3: Agency Subsamples (Full Model Specification)**

| Variable | DoD | HHS | NASA | VA |
|---|---|---|---|---|
| Log exposure | -0.0087 | 0.0014 | -0.0511* | -0.0219 |
| | (0.0170) | (0.0179) | (0.0295) | (0.0366) |
| Log market cap | 0.0330*** | 0.0327*** | 0.0362*** | 0.0276*** |
| | (0.0038) | (0.0026) | (0.0065) | (0.0040) |
| Leverage | 0.0739*** | 0.0523*** | 0.0532* | 0.0141 |
| | (0.0163) | (0.0092) | (0.0318) | (0.0273) |

| | | | | |
|---|---|---|---|---|
| R&D ratio | 0.0082*** | 0.0000** | -0.1739*** | 0.3136** |
| | (0.0017) | (0.0000) | (0.0589) | (0.1296) |
| Log total assets | -0.0323*** | -0.0304*** | -0.0415*** | -0.0329*** |
| | (0.0051) | (0.0042) | (0.0094) | (0.0091) |
| Firm FE | Yes | Yes | Yes | Yes |
| Time FE | Yes | Yes | Yes | Yes |
| Observations | 9,299 | 6,818 | 3,718 | 2,311 |
| $R^2$ (within) | 0.0218 | 0.0291 | 0.0233 | 0.0171 |

Notes: Each column estimates the full M3 specification on the subsample of firms classified under each agency. Firms may appear in multiple subsamples. Standard errors clustered by firm in parentheses. *** p<0.01, ** p<0.05, * p<0.10

## Table 4: Interaction Effects: COVID Period and Political Cycles

| **Variable** | **Table 4a** | **Table 4b** |
|---|---|---|
| | COVID Interaction | Political Cycles |
| Log exposure | 0.0034 | 0.0031 |
| | (0.0149) | (0.0147) |
| Log exposure × COVID | -0.0196 | |
| | (0.0110) | |
| Log exposure × Democrat | | -0.0180 |
| | | (0.0119) |
| Log market cap | 0.0345*** | 0.0343*** |
| | (0.0026) | (0.0026) |
| Leverage | 0.0636*** | 0.0637*** |
| | (0.0104) | (0.0105) |
| R&D ratio | 0.0000** | 0.0000** |
| | (0.0000) | (0.0000) |
| Log total assets | -0.0330*** | -0.0327*** |
| | (0.0037) | (0.0036) |
| Firm FE | Yes | Yes |

| | | |
|---|---|---|
| Time FE | Yes | Yes |
| Observations | 14,104 | 14,104 |
| $R^2$ (within) | 0.0251 | 0.0253 |

Notes: Dependent variable is monthly stock return. Two-way fixed effects panel regression. Standard errors clustered by firm in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

**Table 5: Fama-MacBeth Cross-Sectional Regressions**

| Variable | Mean Coefficient | Std Error | t-statistic |
|---|---|---|---|
| Intercept | 0.0999 | 0.1098 | 0.91 |
| Log exposure | 0.0818 | 0.0697 | 1.17 |
| Log market cap | 0.0127 | 0.0130 | 0.97 |
| Leverage | 0.0995 | 0.0663 | 1.50 |
| R&D ratio | -0.0259 | 0.0350 | -0.74 |
| Log total assets | -0.0129 | 0.0186 | -0.69 |
| Cross-sections (years) | | 11 | |

Notes: Fama-MacBeth procedure. Annual cross-sectional regressions estimated for each year 2014-2024; reported coefficients are time-series averages. Standard errors are the time-series standard deviation of annual coefficients divided by $\sqrt{T}$. No fixed effects (absorbed by cross-sectional variation). Results are directionally consistent with panel estimates but none reach conventional significance, reflecting the small number of time periods (T=11).

## 8.3 Appendix C - AI Disclosure

Artificial intelligence tools, specifically Claude (Anthropic), were used in a transparent manner throughout the preparation of this thesis. AI assistance was employed to support three aspects of the research process: initial identification of entity relationships during the data construction stage, language refinement and clarity improvements during writing, and debugging of Python code during pipeline development. In all cases, the outputs were reviewed critically by the authors, and all analytical decisions, methodological choices, and interpretations are entirely our own.

We were mindful of the risks associated with generative AI, particularly the possibility of inaccurate or hallucinated outputs. Accordingly, AI tools were not used for factual claims, academic arguments, or empirical interpretation. All ideas, reasoning, and conclusions presented in this thesis are original.