

NBER WORKING PAPER SERIES

DO SOCIAL NORMS SUBSTITUTE FOR ENFORCEMENT?
EVIDENCE FROM PUBLIC OFFICIALS' HOME PURCHASES IN SINGAPORE

Tomasz Piskorski
Amit Seru
Jian Zhang
Chun Zhao

Working Paper 35756
<http://www.nber.org/papers/w35756>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2026

We benefited immensely from discussions with Chang-Tai Hsieh, Shengwu Li, Jesse Shapiro, and Luigi Zingales. We also thank Sumit Agarwal, Souphala Chomsisengphet, Stephen Haber, Erik Hurst, Dan Kessler, Pete Klenow, David Laibson, Ross Levine, Jessica Pan, Ivan Png, Wenlan Qian, Raghuram Rajan, Tarun Ramadorai, David Reeb, Paola Sapienza, Bernard Yeung, and seminar participants at Hoover Long Run Prosperity Conference at Koch University, Path Towards Prosperity: Institutions & Development at Hoover Institution, National University of Singapore, Stanford University, and University of Chicago for helpful comments. Some of the survey data analyzed in this paper were collected by the Asian Barometer Project, Waves 2–5. We thank the Asian Barometer Project and its participating institutions for making these data available. The Asian Barometer Project Office is solely responsible for data distribution, and the views expressed herein are the authors' own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Tomasz Piskorski, Amit Seru, Jian Zhang, and Chun Zhao. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Do Social Norms Substitute for Enforcement? Evidence from Public Officials' Home Purchases in Singapore

Tomasz Piskorski, Amit Seru, Jian Zhang, and Chun Zhao

NBER Working Paper No. 35756

September 2026

JEL No. E59, G14, G18, K29

ABSTRACT

We study whether pro-integrity social norms can substitute for formal enforcement in deterring misconduct by public officials. Singapore, widely regarded as one of the world's least corrupt countries, provides a sharp setting to examine whether decades of successful anti-corruption enforcement can produce self-sustaining norms. Using universe-level housing transactions, we identify informed home purchases by civil servants around expansions of the Mass Rapid Transit (MRT) system. Relative to a matched uninformed control, civil servants disproportionately purchase homes near planned—but not yet publicly announced—stations, with the effect concentrated one to two years before public announcements. The behavior is concentrated among mid-level officials and agencies connected to rail planning, generates meaningful private gains, and also appears among relatives, consistent with information leakage. Stronger formal enforcement substantially reduces both direct and indirect misconduct. We combine the transaction evidence with independent survey measures of perceived corruption, public trust, and pro-integrity norms and develop a dynamic model in which norms are persistent but fragile, enforcement both directly deters misconduct and indirectly sustains norms, and policymakers learn about the persistence of norms from noisy integrity signals. Observed pre-tightening misconduct, deterioration in survey-based integrity measures, and the subsequent decline in misconduct following stronger enforcement are more consistent with policymakers learning that norms have limited persistence than with a response to transitory noise. The findings suggest that even in high-integrity societies, social norms may not be sufficiently self-sustaining to maintain a low-corruption equilibrium without continued formal enforcement.

Tomasz Piskorski
Columbia University
and NBER
tp2252@columbia.edu

Jian Zhang
The University of Hong Kong (HKU)
Business School
zhangj1@hku.hk

Amit Seru
Stanford University
Stanford Graduate School of Business
and NBER
aseru@stanford.edu

Chun Zhao
Stanford University
Graduate School of Business
Department of Finance
chunzhao@stanford.edu

I. Introduction

There is a long-standing debate over whether pro-integrity social norms can substitute for formal enforcement in curbing misconduct. In principle, societies with strong norms may rely less on costly monitoring and sanctions, lowering enforcement costs and improving efficiency (e.g., [Guiso et al. 2004, 2006, 2009](#); [Bidner and Francois 2013](#); [Acemoglu and Jackson 2015](#); [Benabou and Tirole 2026](#)). Yet the canonical law-and-economics view emphasizes that credible rules and penalties remain essential to discipline behavior (e.g., [Becker 1968](#); [Djankov et al. 2003](#); [La Porta et al. 1998, 1999](#); [Posner 1997, 2002](#); [Shleifer and Vishny 1998](#)). This trade-off remains highly relevant in practice. Current debates in the United States over reducing regulatory oversight, including in financial markets, are often motivated by the potential to lower monitoring costs and regulatory burdens. Yet weaker enforcement may also undermine deterrence and gradually erode the very norms that sustain lawful behavior.

A particularly informative setting in which to study this question is the misuse of privileged official information for personal gain. Analogous to corporate insider trading, informed trading by government officials can distort markets, redistribute wealth, erode trust in public institutions, and fuel perceptions of unfairness that exacerbate inequality and political instability. High-profile episodes—including scrutiny of U.S. legislators' COVID-era trades, the 2021 resignations of two Federal Reserve regional bank presidents over trading controversies, and the 2025 resignation of a Fed governor amid allegations of insider trading by their relatives—have intensified debate and renewed calls for reform.¹ Yet systematic evidence on the prevalence of such behavior remains limited, making it difficult to determine whether deterioration in the perceived integrity environment reflects actual misuse of privileged information or merely noisy signals.

Singapore provides a unique opportunity to study this question. Corruption was once widespread but declined following decades of sustained anti-corruption enforcement, leaving Singapore widely regarded as one of the world's least corrupt countries. If social norms are ever able to largely substitute for formal enforcement, one would expect this to occur in such a setting. At the same time, Singapore's unusually rich administrative data allow us to construct transaction-level measures of behavior strongly indicative of misconduct rather than relying solely on detected cases or noisy signals about the integrity environment. This makes it possible to distinguish between deterioration in the perceived integrity environment that reflects genuine underlying misconduct and deterioration driven by noisy signals, and, in turn, to draw inferences about whether norms were sufficiently persistent to sustain a high-integrity equilibrium without continued formal enforcement.

¹E.g., see "[As Covid Hit, Washington Officials Traded Stocks With Exquisite Timing](#)" (The Wall Street Journal, Oct. 19, 2022). Also see "[Fed Leaders Eric Rosengren, Robert Kaplan to Resign Following Trading Controversy](#)" (The Wall Street Journal, Sept. 27, 2022), and "[Former Fed Governor Violated Central Bank's Trading Rules, New Disclosures Show](#)" (Washington Post, Nov. 15, 2025).

Our empirical strategy combines three complementary sources of evidence. We first construct a transaction-level measure of misconduct by public officials using universe-level housing transactions. We identify informed home purchases around planned—but not yet publicly announced—Mass Rapid Transit (MRT) stations by linking the universe of private housing transactions to civil-service identities. We then study how this behavior, conduct prohibited by the civil-service code of conduct, changes around an enforcement-tightening episode and complement this evidence with independent survey measures of perceived corruption, public trust, and pro-integrity norms. Together, these complementary data provide the empirical foundation for the analysis that follows.

To organize the empirical analysis, we develop a simple dynamic model in which social norms are persistent but fragile: norms carry over from one period to the next but gradually erode when misconduct rises. Governments do not directly observe underlying misconduct and therefore cannot know whether norms remain sufficiently persistent to sustain compliance. Instead, they infer the state of the integrity environment from noisy signals and decide whether continued formal enforcement is warranted. Formal enforcement deters misconduct directly and sustains norms indirectly by increasing the value of maintaining a clean reputation. We map the model to Singapore’s historical trajectory to interpret the interaction between social norms, formal enforcement, and misconduct.

The empirical analysis is organized around two complementary questions implied by the model. The first concerns the presence and distribution of misconduct. If misconduct is present, where in the bureaucracy does it arise, and which of the model’s alternative regimes best characterizes its distribution? Our transaction-level measures allow us to document where informed purchases occur, quantify their consequences, and identify the regime most consistent with the data. The second question concerns whether continued formal enforcement remained necessary to sustain a high-integrity equilibrium. By combining evidence on misconduct before the enforcement tightening, its response afterward, and changes in the broader integrity environment, we assess whether adverse integrity signals reflected genuine underlying misconduct or merely noise. This allows us to draw inferences about whether norms were sufficiently persistent to sustain compliance without continued formal enforcement.

We begin by examining whether public officials engage in informed home purchases before planned rail stations are publicly announced. We assemble the universe of private-property transactions from legal filings in Singapore and link buyers to the Singapore Government Directory Interactive (SGDI) to identify civil servants and trace their transaction histories. To proxy for access to privileged location information, we focus on purchases near planned – but not yet publicly announced – Mass Rapid Transit (MRT) stations around announcement dates. Because properties close to new stations are expected to command positive proximity premia, such pre-announcement purchases are consistent with informed trading. To mitigate confounding determinants of location choice, we construct a matched sample of civil servants and observationally similar non-civil servants as

uninformed controls. A key advantage of our setting is coverage: unlike studies that rely on exposed cases of misconduct, we are able to identify such behavior across the full distribution of officials.

We find clear evidence of informed trading by civil servants: purchase intensity near planned MRT stations is higher for civil servants than for matched controls, with the difference concentrated one to two years before public announcements. Robustness checks support this finding. First, a permutation test, randomly reassigning the civil-servant indicator, produces a tightly centered null, and our estimated effect remains statistically significant based on bootstrapped p-values. Second, the results are insensitive to alternative distance thresholds used to define "nearby" properties. Third, in a placebo exercise using Singapore's highly regulated subsidized public-housing market, where eligibility and resale constraints make informed trading more difficult, we find no differential pre-announcement buying by civil servants relative to matched controls for public-housing units near planned stations.

We next examine the presence and distribution of misconduct to address the first question posed by the model. Using LLM-assisted classification to tag civil-service ranks, and validating the classifications on a hand-checked sample, we find that informed purchases are concentrated among mid-level officials. Consistent with the model, this pattern indicates that the middle-concentration regime best characterizes the data: mid-level officials face the highest net temptation because they have sufficient resources to act while bearing lower expected scrutiny and reputational costs than senior officials, whereas junior officials face binding resource constraints. We also examine proximity to information sources and find that informed purchases are more prevalent among officials employed in agencies involved in rail-line planning, consistent with greater access to nonpublic location information.

A natural alternative explanation is that civil servants are simply better at processing public information rather than relying on privileged information. To examine this possibility, we compare their behavior with two groups possessing high information-processing sophistication—real-estate agents and corporate board members—using the same event window and matching strategy. Neither group exhibits elevated pre-announcement purchases near planned stations relative to matched controls. This absence of similar behavior among sophisticated private actors strengthens the interpretation that civil servants' purchases reflect access to nonpublic information rather than superior processing of public information.

We next examine the economic consequences of informed trading by public officials, beginning with direct benefits to the officials themselves and an indirect spillover channel through their relatives. Because housing provides both consumption services and investment returns, measuring the gains from informed trading is not straightforward. We therefore focus on realized sales and benchmark-adjusted returns. Two findings emerge. First, civil servants do not pay price premia at the time of

purchase. Second, among properties that are subsequently sold, informed purchases generate positive benchmark-adjusted returns over shorter holding periods than comparable transactions, consistent with an investment motive. These gains are economically meaningful relative to civil-service incomes. We also document a spillover channel: relatives of civil servants engage in informed purchases, earn excess benchmark-adjusted returns, and, on average, exit their investments earlier than matched controls. Together, these findings indicate that public officials and their relatives derive substantial private financial benefits from informed home purchases.

We then examine the broader implications of misconduct. First, informed trading by public officials generates wealth transfers from uninformed sellers to informed buyers with similar incomes, implying welfare losses that are unlikely to be offset by redistribution from wealthier to poorer households. Second, we consider a stronger mechanism in which officials purchase properties early and subsequently influence station placement. We find no evidence consistent with such influence, which accords with our heterogeneity results showing that informed purchases are concentrated among mid-level officials who are unlikely to possess such authority.

Finally, we turn to the second question posed by the model: whether continued formal enforcement remained necessary to sustain a high-integrity equilibrium. We study the enforcement tightening, including the Securities and Futures (Amendment) Act 2012 and the broader strengthening of anti-corruption enforcement, to examine how misconduct responds to stronger formal enforcement.² Using time-series evidence around future-station announcements, we show that prior to the enforcement tightening not only civil servants but also their relatives increased purchase activity and earned significant abnormal returns, consistent with information leakage through family networks. Following the enforcement tightening, informed purchases by officials disappear, with no offsetting substitution toward relatives, indicating that stronger formal enforcement shut down both the direct and indirect channels of misconduct.

We then combine this evidence with independent survey measures of perceived corruption, public trust, and pro-integrity norms to interpret the enforcement episode. The survey evidence documents deterioration in the broader integrity environment around the same period but, on its own, cannot distinguish real behavioral deterioration from changing perceptions or other noise. Our transaction-level evidence provides the missing information: meaningful misconduct was present before the enforcement tightening and declined broadly afterward. Taken together, the evidence is more consistent with policymakers responding to genuine weakening in the integrity environment than with reacting to noisy signals alone. This, in turn, suggests that even after

²See [Quah \(2020\)](#) for detailed discussions. As noted therein, “even though the number of corruption-related complaints received by the CPIB declined during 2006–2011, several high-profile corruption scandals were uncovered in 2011 and 2012.” Most of these cases came to public attention and triggered reviews and enforcement tightening in 2012, as underscored by [Lee \(2012\)](#).

decades of enforcement, social norms may not be sufficiently persistent to sustain a high-integrity equilibrium, even in societies long regarded as having low levels of corruption and strong ethical norms. Continued formal enforcement may therefore remain necessary to preserve compliance.

Our paper connects to classic law-and-economics work on how social norms and formal enforcement deter misconduct by public officials (e.g., [Acemoglu and Verdier 2000](#); [Becker 1968](#); [La Porta et al. 1999](#); [Posner 1997, 2002](#); [Shleifer and Vishny 1998](#)) and to the broader law-and-finance tradition on how legal quality and enforcement shape behavior and market outcomes ([La Porta et al. 1998](#); [Shleifer and Vishny 1993, 1997](#); [Rajan and Zingales 1998, 2003](#)). It also relates to models with multiple compliance equilibria and the evolution of norms and cooperation ([Kandori 1992](#); [Tirole 1996](#); [Tabellini 2008](#); [Bidner and Francois 2013](#); [Acemoglu and Jackson 2015, 2017](#); [Bisin and Verdier 2024](#); [Benabou and Tirole 2026](#)). Empirically, [Fisman and Miguel \(2007\)](#) show that diplomats' parking violations vary with home-country corruption norms and fall sharply when local legal enforcement strengthens, while [Lane et al. \(2023\)](#) show experimentally that law can shape social norms. Complementing this literature, our evidence from Singapore suggests that even when decades of stringent enforcement have helped establish a low-corruption equilibrium, the resulting norms may not be self-sustaining in the absence of continued formal enforcement. Maintaining probity may therefore require enforcement to remain both credible and salient.

Our paper also contributes to empirical work on public-sector corruption and anti-corruption enforcement. Studies exploiting anti-corruption crackdowns, audits, monitoring, and judicial sanctions document the effects of formal oversight, with substantial variation in persistence and across activities and public officials ([Olken 2007](#); [Bobonis et al. 2016](#); [Avis et al. 2018](#); [Colonnelli and Prem 2022](#); [Axbard 2025](#)). Related work uses administrative data to uncover otherwise concealed private benefits of public office ([Agarwal et al. 2020](#)). We extend this work to a high-integrity bureaucracy and to a hard-to-detect form of public-office misuse: the use of nonpublic government information for private gain. Linking the universe of housing transactions to the identities of individual civil servants allows us to construct a transaction-level behavioral proxy for misconduct that does not depend on cases being formally exposed and to examine how it responds to stronger enforcement. Combining this proxy with survey-based measures of the broader integrity environment allows us to assess whether changes in perceived corruption coincide with changes in underlying misconduct ([Olken 2009](#)).

Beyond this literature, our paper is related to empirical work that uses economic structure and unusual data to uncover hidden behavior and latent institutional forces. In closely related recent work, [Andirin et al. \(2025\)](#) develop a model showing that preventive repression makes protests less predictable and use novel security-alert data containing advance warnings of protests to construct a measure of this otherwise difficult-to-observe form of state control. Although our question is different, our approach is related. In our setting, survey evidence captures noisy

signals of deteriorating integrity available to policymakers, while administrative transaction data reveal underlying misconduct. Together with the response of misconduct to stronger enforcement, our model helps distinguish genuine deterioration from fluctuations in perceptions and separate compliance sustained by pro-integrity norms from compliance produced by deterrence.

Finally, we contribute to work on informed trading by politicians and government personnel. Studies report mixed performance for U.S. legislators. Some find that mimicking their portfolios yields annual outperformance of 6–12% (Ziobrowski et al. 2004, 2011), whereas others find no outperformance (Eggers and Hainmueller 2013; Belmont et al. 2022). Related work suggests that connected brokerages may provide politicians with stock tips in exchange for political access or information (Stephan et al. 2021). More broadly, we relate to research on private information and misconduct in financial and political intermediation, including analyst tipping (Irvine et al. 2007), illegal insider-trading networks (Ahern 2017), early access to news feeds and SEC filings (Hu et al. 2017; Rogers et al. 2017), sovereign-rating news (Michaelides et al. 2015), regulatory information leakage (Reeb et al. 2014), lobbying (Bertrand et al. 2014, 2020), mortgage-backed-securities misreporting (Piskorski et al. 2015; Griffin and Maturana 2016), and financial-advisor misconduct (Egan et al. 2017). We add to this literature by using universe-level data to study career civil servants' housing purchases around infrastructure announcements, allowing us to examine heterogeneous information access, potential leakage to relatives, realized gains, and the role of regulation.

The remainder of the paper is organized as follows. Section II lays out the theoretical framework and testable implications. Section III describes the institutional background, the data, and the methodology. Section IV presents the main results of our analysis. Section V presents results on welfare effects of informed trading. Section VI presents evidence on the role of regulation in curbing informed trading by government officials. Section VII concludes.

II. A Model of Norms, Enforcement, and Learning

To guide our empirical analysis, we develop a simple, multi-period framework in the spirit of Tirole (1996), Acemoglu and Jackson (2015, 2017), and Benabou and Tirole (2026), in which formal enforcement and social norms jointly discipline misconduct by government officials. Within each rank, officials differ in moral costs and face three forces: a private payoff from corruption, the risk of detection with penalties, and a reputational payoff that accrues to those who keep a clean record. Because monitoring is imperfect, a clean record is only an imperfect signal of integrity, and the social value of reputation rises with the overall climate of honesty. Enforcement therefore operates through two channels: it directly raises expected legal costs and it indirectly increases the informational content—and thus the value—of a clean record by sustaining a higher honesty norm. This complementarity implies that stronger enforcement may both deter misconduct and deepen the reputational premium for honesty.

We allow the government to choose enforcement over time while learning about the persistence of norms, ϕ , which governs how current honesty carries forward into future reputational capital. The government does not observe true misconduct; instead, it infers behavior from noisy signals of the public reputational stock, and it updates its belief about ϕ period by period.

We apply the framework to interpret Singapore’s post-1965 trajectory. We read the anti-corruption push throughout the second half of the twentieth century as an endogenous response to low initial norms, undertaken to build state capacity, secure public trust, and support investment-led growth. In the model, tighter enforcement both depresses perceived misconduct and raises the honesty norm. This channel corresponds to Singapore’s widely recognized experience, where sustained enforcement helped entrench a culture of integrity. The government learns about norm persistence, ϕ , from noisy signals and adjusts enforcement as beliefs evolve. By the late 2000s, adverse signals may have indicated deterioration in the integrity environment. The model clarifies when a tightening in response to such signals reflects noise and when it reflects learning about limited norm persistence.

The framework delivers two empirical implications. First, misconduct need not be uniform across the bureaucracy. Depending on illicit returns, expected sanctions, and reputational costs, misconduct can concentrate either at the extremes of the hierarchy or among mid-level officials. Observing where misconduct appears therefore helps identify the relevant incentive regime and the ranks where enforcement is most valuable.

Second, the framework distinguishes two interpretations of an enforcement tightening following adverse integrity signals. If the adverse signal is noise, underlying misconduct should be low or stable before the tightening. If the signal reflects learning about limited norm persistence, underlying misconduct should already be elevated and should decline when enforcement strengthens. Unlike policymakers in the model, we construct transaction-level measures of misconduct, allowing us to distinguish between these mechanisms empirically.

II.A. Model Setup

II.A.1. Environment and Agents

Time is indexed by periods $t = 0, 1, 2, \dots$. In each period, a new cohort of officials enters. Officials are partitioned into three ranks $i \in \{L, M, H\}$ (low, middle, high) with population shares $\alpha_i > 0$ and $\sum_i \alpha_i = 1$. Within each rank, officials are heterogeneous in their intrinsic moral cost of misconduct $\theta \in [0, 1]$, drawn i.i.d. across officials and periods from a continuous distribution (we will use the uniform distribution in applications). A higher θ represents a larger psychic/moral disutility from committing a corrupt act. Short-lived officials choose whether to engage in misconduct in order to maximize their personal utility.

The key state variable is the reputational capital, R_t , which captures the social value of having a clean record. At the end of each period, society experiences ongoing structural transitions, possibly triggered by cultural or technological shocks. As a result, the social norm evolves according to:

$$(1) \quad R_t = \phi R_{t-1} + (1 - \phi)\rho x_t + v_t, \quad x_t \equiv \sum_{i \in \{L, M, H\}} \alpha_i x_{i,t},$$

where $\phi \in [0, 1]$ is the (unknown) persistence of norms, $\rho > 0$ is a scale parameter, $x_{i,t} \in [0, 1]$ is the honesty share in rank i during period t , and v_t is a random shock. A benevolent government seeks to maintain the integrity capital R and improve the contemporaneous social honesty rate x , equivalently reducing corruption.

II.A.2. Officials' Payoffs and Rank-Specific Parameters

An official in rank i who commits a corrupt act receives a private benefit $B_i > 0$ but faces a formal penalty $F_i > 0$ if detected with probability p_t . Behavior is also disciplined by reputation. At the start of period t , officials take the public reputational capital, $\hat{R}_{t-1}$, as given. This capital reflects a social consensus on the value of integrity, and its tangible return is weighted by a rank-specific parameter $\gamma_i \geq 0$. We allow for a period- t common disturbance to behavior, ϵ_t , that officials observe when choosing actions but the government does not. For an official with intrinsic moral cost θ , the utilities for honesty and corruption are:

$$(2) \quad U_{i,t}^{\text{honest}} = \gamma_i \hat{R}_{t-1},$$

$$(3) \quad U_{i,t}^{\text{corrupt}} = B_i - \theta - p_t F_i + (1 - p_t) \gamma_i \hat{R}_{t-1} - \epsilon_t.$$

Corrupt officials who escape detection with probability $1 - p_t$ maintain their clean record and thus capture the same reputational payoff as honest officials. The term $(1 - p_t) \gamma_i \hat{R}_{t-1}$ captures a key aspect of this incentive structure: the moral hazard of free-riding. In a high-integrity society, a clean public record is a valuable signal of incorruptibility. Dishonest officials who are not caught can exploit this by "free-riding" on the high reputational capital, enjoying its benefits without having earned them. This moral hazard directly disincentivizes genuine honesty and, over time, erodes the value of reputation (R). Enforcement therefore serves to curb this free-riding, strengthening the social norm.

The within-rank honesty share, $x_{i,t} \in [0, 1]$, is the measure of types for whom $U_{i,t}^{\text{honest}} \geq U_{i,t}^{\text{corrupt}}$, and aggregate honesty is $x_t = \sum_i \alpha_i x_{i,t}$. The government does not observe the specific parameters for each rank and only perceives the aggregate characteristics of the official population. We define the observable aggregate parameters for the government's problem as $\gamma \equiv \sum_i \alpha_i \gamma_i$, $F \equiv \sum_i \alpha_i F_i$, and $B \equiv \sum_i \alpha_i B_i$.

II.A.3. Government Policy and Objective

At the start of each period t , the government acts as a benevolent social planner, choosing a uniform enforcement intensity p_t to maximize social norms.³ Enforcing the law is costly, described by an increasing and convex cost function $C(\cdot)$, such that $C(p) = \frac{\kappa}{2}p^2$. The government's objective is to balance the social benefit from high reputational capital against the costs of enforcement. Since the true norm persistence ϕ is unknown and officials' actions have not yet occurred, the government maximizes the expected value of its per-period objective function based on its available information, $\mathcal{I}_{t,\text{pre}}^G$. The optimization problem is:

$$(4) \quad \max_{p_t \in [0,1]} \alpha \mathbb{E}[R_t | \mathcal{I}_{t,\text{pre}}^G] - C(p_t),$$

where α captures the relative weight the government places on maintaining a high-integrity public sector, summarized by the norm state R_t . Without loss of generality we normalize α to 1. For simplicity, we model the government choosing enforcement to maximize the expected end-of-period norm state, R_t , net of current enforcement costs. Since R_t is also the norm stock carried into period $t + 1$, this formulation captures the essential trade-off between immediate enforcement costs and the expected reputational outcome of its policy choice.

II.A.4. Information and Timing

Information and Beliefs. The model's informational environment is defined by the true, unobservable (latent) reputational capital, R_t , starting from an initial state, R_0 . The government cannot observe R_t directly. Instead, at the end of each period $t \geq 0$, it receives a noisy public signal of the reputational capital, $\tilde{R}_t \equiv R_t + \eta_t$, where η_t is a random noise. The government's decision-making is guided by its belief about the persistence of norms, $\hat{\phi}_t$. This belief, along with the known initial state and the history of past signals, constitutes the government's information set, $\mathcal{I}_{t,\text{pre}}^G = \{\hat{\phi}_t, \text{history of } (\tilde{R}_\tau)_{\tau \leq t-1}\}$. The central friction is that the government is informationally constrained: it knows neither the true persistence parameter ϕ nor the contemporaneous honesty rates $(x_{i,t})_i$. Officials share the government's information set and, crucially, observe the enforcement policy p_t before choosing their actions.

Sequence of Events. The timing within each period t unfolds as follows:

- (1) *Government Action:* Based on its information set $\mathcal{I}_{t,\text{pre}}^G$, the government chooses the enforcement intensity p_t to maximize its objective function.
- (2) *Officials' Reaction:* Observing p_t and ϵ_t , officials decide whether to act honestly or corruptly. These actions collectively determine the aggregate honesty share for the period, x_t .

³A uniform policy is optimal given the government's information constraints. Lacking knowledge of rank-specific parameters, the government cannot design a targeted strategy and thus sets a single enforcement intensity for all officials, i.e., $p_{L,t} = p_{M,t} = p_{H,t} = p_t$.

- (3) *State Evolution and Signal Realization*: The latent reputational capital evolves according to its law of motion, $R_t = \phi R_{t-1} + (1 - \phi) \rho x_t + v_t$. A new public signal, $\tilde{R}_t = R_t + \eta_t$, is then realized.
- (4) *Belief Updating*: The government compares the realized signal $\tilde{R}_t$ to its prior forecast. Using this prediction error, it updates its belief about norm persistence, revising it from $\hat{\phi}_t$ to $\hat{\phi}_{t+1}$. This step concludes period t and establishes the initial belief for period $t + 1$.

This sequence establishes a sequential problem where the government must learn about the underlying social structure while simultaneously trying to control behavior through its policy choices.

II.B. Equilibrium and Testable Implications

II.B.1. Equilibrium and Comparative Statics

Let the public state at the start of period t be the government's belief about persistence, $\hat{\phi}_t$, and the prior period's reputation, $\hat{R}_{t-1}$. A period- t equilibrium consists of an enforcement policy $p_t^* \in [0, 1]$, officials' anticipated honesty shares $(x_{i,t})_i$, aggregate honesty x_t , and a resulting reputational stock R_t such that: (i) given p_t^* , each official's choice is optimal, determining $x_{i,t}$; and (ii) the government's choice of p_t^* maximizes its subjective expected welfare based on its beliefs. In Appendix E, we solve the model and prove all statements below. Within the interior region where $0 < x_{i,t} < 1$, the honesty share $x_{i,t}$ is strictly increasing in enforcement p_t and the level of reputation $\hat{R}_{t-1}$. Below, we discuss several comparative statics that directly inform our empirical analysis.

Complementarity between Social Norms and Legal Enforcement. Enforcement and norms are complements. Stricter monitoring (p_t) is more effective in a high-reputation environment ($\hat{R}_{t-1}$) because the potential loss from being caught—comprising both formal penalties and the forfeiture of a valuable clean record—is larger. This raises the marginal return to enforcement. Formally, for any interior rank, the cross-partial derivative is positive: $\frac{\partial^2 x_{i,t}}{\partial \hat{R}_{t-1} \partial p_t} > 0$. Therefore, due to the complementarity discussed above, a higher $\hat{R}_{t-1}$ makes enforcement more potent, increasing the marginal benefit of each unit of enforcement effort. Formally, for interior equilibrium, the optimal enforcement increases with social norm: $\frac{\partial p_t^*}{\partial \hat{R}_{t-1}} > 0$.

High social norms lower the enforcement floor. A direct policy implication concerns the enforcement floor, $\bar{p}_t$, needed to sustain zero corruption (corner equilibrium). $\bar{p}_t$ is decreasing in $\hat{R}_{t-1}$, $\frac{\partial \bar{p}_t}{\partial \hat{R}_{t-1}} < 0$. A society with a larger reputational stock requires a lower monitoring intensity to maintain full compliance in that rank. In equilibrium, the government's endogenous choice of enforcement, p_t , balances these falling thresholds against the costs of enforcement. Thus, a high initial investment in enforcement can build sufficient reputational capital, $\hat{R}$, to allow for lower long-term enforcement without compromising near-zero corruption.

Norms Require Sustained Enforcement to Avert Decline. The model features a feedback loop. An unexpected lapse in enforcement (p_t falls below optimal) immediately reduces current honesty ($x_t \downarrow$).

This, through the law of motion, leads to a lower-than-expected true reputation stock ($R_t \downarrow$). This initial damage then propagates into the future through two distinct channels. The lower reputation stock directly weakens future incentives, as the total stakes of misconduct $F + \gamma \hat{R}_t$ are now lower for all officials. Concurrently, the low R_t generates a negative signal for the government, which, through its learning process, leads to more pessimistic beliefs about the persistence of norms ($\hat{\phi}_{t+1} \downarrow$). Consequently, in the next period, the government faces a more fragile situation—weaker underlying incentives and a more pessimistic outlook—which can create a vicious cycle that is costly to reverse. This dynamic highlights that maintaining a high-norm equilibrium requires consistent and sustained enforcement.

II.B.2. Connecting Model and Empirics: Singapore and Testable Implications

To conclude our theoretical analysis, we map the framework to Singapore’s historical trajectory and motivate the empirical analysis that follows. As described in the institutional background that follows in Section III, Singapore transitioned from a high-corruption environment in the 1950s and 1960s to one widely regarded as having exceptionally strong public-sector integrity following decades of sustained anti-corruption efforts after independence in 1965. In the model, this corresponds to lowering the unit cost of enforcement (a smaller κ) and choosing a high initial enforcement p_0 when the reputational stock is perceived low (baseline signal $\tilde{R}_0$) and norm persistence ϕ is unknown. Alternatively, the early post-independence period can also be interpreted as one in which policymakers placed a higher value on building public-sector integrity (a higher α in the government’s objective), making strong anti-corruption enforcement an explicit policy priority. Early and sustained crackdowns, pay/process reforms, and visible prosecutions are captured by persistently higher $\{p_t\}$, which both directly deters misconduct and indirectly strengthens the honesty norm R_t by increasing the value of a clean record. As $\tilde{R}_t$ arrives over time, the government updates beliefs about ϕ and adjusts p_t with the posterior. When entering norms are high, the comparative statics derived above imply a lower minimum enforcement floor $\bar{p}_t = B/(F + \gamma \hat{R}_{t-1})$ and a lower interior policy $p_t^{\text{in}*} = (\rho/\kappa) \mathbb{E}_t[(1 - \hat{\phi}_t)(F + \gamma \hat{R}_{t-1})]$ as perceived persistence $\hat{\phi}_t$ rises.

This mapping closely matches Singapore’s historical experience, where decades of strict and credible enforcement helped entrench a culture of integrity. The central question is what happens once such a high-integrity equilibrium has been established. Can governments safely rely on norms to sustain compliance, or does maintaining such an equilibrium continue to require active formal enforcement? The difficulty is that governments cannot answer this question directly. Policymakers do not fully observe underlying misconduct and consequently cannot directly assess whether norms remain sufficiently persistent to sustain compliance. Instead, they must infer the state of the integrity environment from noisy signals and decide whether continued formal enforcement is warranted.

Our empirical setting is unusual because, unlike the government in the model, we construct a transaction-level measure of misconduct from administrative data that does not rely on detected

or prosecuted cases. This provides direct evidence on the underlying behavior that policymakers themselves cannot observe. This additional information allows us to address two questions that are difficult to answer using the noisy signals available to governments alone.

The first question concerns the presence and cross-sectional distribution of misconduct. If misconduct is present, which of the model's alternative regimes best characterizes its distribution across the bureaucracy? Because governments do not directly observe misconduct, establishing whether it exists, and where it is concentrated, is informative both about the prevalence of misconduct and about where formal enforcement is likely to be most valuable. Sections [IV](#) and [V](#) address this question by introducing our transaction-level measure of misconduct, documenting where misconduct occurs within the bureaucracy, and quantifying its economic consequences.

The second question concerns whether the observed misconduct implies that continued formal enforcement remains necessary to sustain a high-integrity equilibrium. Section [VI](#) studies this question through the 2011–2012 enforcement tightening. We first examine how misconduct responds to stronger formal enforcement and then characterize the broader integrity environment using independent evidence on perceived corruption, public trust, and pro-integrity norms. By combining them with our transaction-level measure of misconduct, we investigate whether meaningful misconduct was already present before the enforcement tightening and whether it subsequently responded to stronger enforcement. Taken together, these pieces of evidence allow us to assess whether the deterioration in the integrity environment reflected a genuine weakening of compliance that policymakers correctly inferred, or merely noisy signals that did not warrant a strengthening of formal enforcement.

The two questions implied by the model can now be stated more formally. Formal derivations are provided in Appendices [E2](#) and [E3](#).

The Presence and Distribution of Misconduct (Question 1). The model does not necessarily predict that misconduct will arise. If social norms are sufficiently strong or enforcement is sufficiently stringent, all officials may choose honest behavior and no misconduct is observed. When misconduct does arise, however, the model predicts that its concentration across the bureaucracy depends on the relative importance of illicit returns, formal sanctions, and reputational costs. In particular, the model admits two possible regimes:

- (1) *Extreme-concentration regime:* When illicit returns are sufficiently high relative to sanctions and reputational costs, misconduct concentrates at the extremes (high and low ranks), with the middle rank exhibiting more honesty due to balanced incentives.
- (2) *Middle-concentration regime:* When the extremes are effectively cornered—either heavily deterred (high rank) or lacking sufficient returns (low rank)—only the middle rank remains at an interior solution responsive to enforcement changes.

These two regimes imply fundamentally different enforcement priorities. The extreme-concentration regime requires a two-pronged strategy that addresses both ends of the hierarchy, perhaps with a particular focus on the high-stakes misconduct of senior officials. In contrast, the middle-concentration regime suggests that enforcement should be aimed squarely at mid-level bureaucrats. However, the government's inability to observe rank-specific behavior creates a risk of misallocating enforcement resources. Specifically, it may expend costly monitoring on ranks that are already fully compliant while failing to target those where misconduct is concentrated.

The Persistence of Norms and the Role of Formal Enforcement (Question 2). The model asks whether a high-integrity equilibrium, once established, becomes largely self-sustaining because norms are sufficiently persistent, or whether it continues to require active formal enforcement. Because governments observe only noisy signals about the integrity environment, they cannot answer this question directly. Tightening enforcement is generally triggered by negative signals about reputation ($\tilde{R}_t < \mathbb{E}_t[\tilde{R}_t]$), which can arise from two distinct mechanisms:⁴

- (1) *Noise-driven overreaction:* If $\hat{\phi}_t = \phi$ (correct prior) and $\eta_t < 0$ (negative noise), the government temporarily overreacts. Under this scenario, pre-signal corruption is low and stable, enforcement tightening is transitory, and corruption exhibits mean reversion or remains constantly near-zero.
- (2) *Correct learning about low persistence:* A negative signal may instead reveal that the integrity environment is weaker than the government had believed. This occurs when the government overestimates norm persistence ($\hat{\phi}_t > \phi$) and the signal is informative ($\eta_t \approx 0$). Before the signal, enforcement is too weak because the government expects norms to discipline behavior on their own. Corruption should therefore be relatively high, or rising, before the policy response. Once the signal arrives, the government revises its belief downward, leading to a sustained tightening of enforcement and a lasting reduction in misconduct.

The key empirical question therefore concerns whether norms are sufficiently persistent for a high-integrity equilibrium to be sustained without continued formal enforcement. If norms are sufficiently persistent, some deterioration in the perceived integrity environment should not be accompanied by meaningful underlying misconduct, and any subsequent strengthening of formal enforcement would most naturally reflect a response to noisy signals rather than genuine behavioral deterioration. By contrast, if deterioration in the perceived integrity environment is accompanied by elevated underlying misconduct, stronger formal enforcement is consistent with policymakers correctly learning that norms are less persistent than previously believed and that continued enforcement remains necessary to sustain compliance.

⁴Assuming ϵ_t is small.

III. Institutional Background, Data, and Methodology

III.A. Institutional Background

III.A.1. The Country and Government

Singapore is an island-state with a total land area of 718.3 square kilometers (km^2) and a population size of 5.9 million as of 2023 (Source: Department of Statistics, Singapore). Because usable land is scarce, high-rise, high-density living is the norm in Singapore, one of the most densely populated countries in the world.

Singapore historically faced significant challenges with corruption. However, following independence in 1965, the government—led by Lee Kuan Yew and the People’s Action Party (PAP)—implemented rigorous anti-corruption measures. These included the strict enforcement of comprehensive laws, supported by institutions such as the Corrupt Practices Investigation Bureau (CPIB), alongside efforts to cultivate a culture of integrity. As a result, corruption has been substantially reduced over the ensuing decades. Today, Singapore is widely regarded as one of the least corrupt countries in the world, consistently ranking highly on transparency and governance indices, including Transparency International’s Corruption Perceptions Index.

In Singapore, the public service includes approximately 141,000 officers across 16 ministries and more than 50 statutory boards. There are three core values upheld by the Singapore Public Service, which include incorruptibility, meritocracy, and pragmatism (Teo 2001). Civil-service salaries, particularly at senior levels, are designed to be competitive with private-sector opportunities in order to attract and retain talent. The Singapore Government has strived to keep its public service lean but efficient by not overly expanding the civil service workforce. Although the Singapore government is the country’s largest employer, representing 4% of the labor force,⁵ the share of Singapore’s workforce employed by the government is comparatively lower than that of many Organisation for Economic Co-operation and Development (OECD) countries.⁶ Singapore’s civil service pay is among the highest in the world, which is aimed at attracting top talent and deterring corruption among high-level officials.⁷

Statutory boards are established and entrusted with autonomous power and rights to carry out selected functions under the supervision of the Ministries. Statutory boards, such as Land Transport Authority (LTA), Housing and Development Board (HDB), and the Central Provident Fund (Board) either directly recruit their own staff or appoint civil servants seconded from the Ministries. They are governed by independent boards of directors made up of representatives from the private sector,

⁵Based on "Expected Increase in Number of Civil Servants" (Ministry of Finance, Singapore, Jan. 19, 2015).

⁶For instance, Japan stands at 6.7%, Germany at 9.6%, the United States at 14.6%, and the United Kingdom at 17.4% (OECD, 2013).

⁷Please see Appendix C for details on Civil Service Pay Grades.

academics, senior officials from public services and trade unions. Similar to other countries, public policies in Singapore such as those related to housing, healthcare, transportation, and education, can wield substantial influence over both economic growth and social welfare.⁸

III.A.2. Regulation of the Public Sector and Prevention of Information Misuse

Singapore is widely regarded as one of the most efficiently run governments, maintaining high standards of governance and accountability within its civil service. To support this, it has implemented multiple layers of institutional regulatory and enforcement frameworks (Khoo 2021). Central to these efforts are two key laws targeting corruption: the Prevention of Corruption Act, enacted in 1960, and the Corruption, Drug Trafficking and Other Serious Crimes (Confiscation of Benefits) Act (CDSA), established in 1992. Both statutes address bribery offenses across the public and private sectors. Complementing these laws, the Corrupt Practices Investigation Bureau (CPIB), an independent agency, plays a crucial role in enforcing anti-corruption measures.

In addition, the Singapore Public Service adheres to a Code of Conduct, which outlines the high standards of behavior expected from public officers, emphasizing principles of integrity, incorruptibility, and transparency. The legal framework was first anchored in the [Prevention of Corruption Act](#) (POCA), enacted in 1960, where Section 6(a) explicitly prohibits any agent (including public officers) from corruptly accepting gratification “as an inducement or reward for showing favor... in relation to his principal’s affairs.” While the statute does not explicitly address the housing-market setting we study, public officers are prohibited from using official information or official position to advance private interests. In response to a public inquiry regarding procurement officers, a director in the Ministry of Finance explicitly stated that “a public officer may not, directly or indirectly, make use of any official information or his official position to further anyone’s private interest, including his own.”⁹ This prohibition is central to our interpretation of pre-announcement purchases by civil servants as misconduct.

Insider trading, in a broader context, is explicitly regulated under the Securities and Futures Act (SFA), which prohibits any individual—including public officers—from trading securities based on non-public, price-sensitive information. The Securities and Futures (Amendment) Act of 2012 expanded insider trading provisions,¹⁰ enhanced the investigative powers of the Monetary

⁸Recent examples of such policies include, among many others, the home-school distance priority for the primary schools’ admissions (Agarwal et al. 2016); the electronic road pricing (ERP) scheme that charges private car owners for the entry into the cordon central business district (Agarwal et al. 2015); the use of certificates of entitlements (COEs) as a market-based mechanism to regulate the Vehicle Quota System (VQS); the introduction of the compulsory saving through the Central Provident Fund (CPF) to encourage saving for retirement and also for housing purchase.

⁹See “Public Officers Governed By Code of Conduct and Discipline” (Ministry of Finance, Singapore, Aug. 27, 2012).

¹⁰Part XII (Market Conduct) of the SFA outlines insider trading provisions and MAS enforcement powers.

Authority of Singapore (MAS), increased penalties for market misconduct, and established a whistle-blowing framework. More broadly, the amendments formed part of a strengthening of Singapore's enforcement framework against misuse of privileged information and market misconduct.¹¹

We use this enforcement-tightening episode, together with the broader strengthening of anti-corruption enforcement around the same period, to examine whether stronger formal enforcement deters the misuse of privileged information by public officials.

III.A.3. Public Transport Policies and Their Announcements

We focus on public-transport infrastructure announcements to study whether access to privileged location information affects public officials' housing transactions. The Land Transport Authority (LTA), a statutory board under the Ministry of Transport, is tasked with overseeing and coordinating public transport planning and operations in Singapore, encompassing taxis, buses, and the Mass Rapid Transit (MRT) rail system. For the MRT operational model, the government via LTA is responsible for major infrastructure provisions, which include the planning and construction of the MRT rail lines and integration of the MRT system with other public transport options.

Panel (a) of Figure 1 displays the comprehensive network of MRT lines in Singapore, including both existing and proposed lines, as of March 2020, while panel (b) shows the evolution of new rail stations. In response to the changing needs of Singapore's expanding population and urban environment, new railway lines such as the Thomson-East Coast Line (TEL), the Jurong Region Line (JRL), and the Cross Island Line (CRL) have been announced. These lines are intended to link residential areas, business districts, and key amenities, thereby expanding the MRT rail network from 178 kilometers (km) to 360 km by 2030.¹²

We focus on public officials' home purchases around announcements of proposed new rail lines. We search all public sources including news releases by the LTA to identify the earliest date on which the new MRT stations are revealed.¹³ We focus on the announcement of new MRT station locations ranging from 1999 to 2019 prior to the COVID outbreak. The details of all new MRT stations and their postal codes that are used in our GIS spatial mapping processes are discussed in Section III.C. The key identifying feature is that station locations are valuable information before public announcement but become public information afterward.

¹¹See "The Securities and Futures (Amendment) Bill 2012" (Ministry of Finance, Singapore, Nov. 15, 2012).

¹²See "Expanding Singapore's MRT: Network to span over 360 km by 2030" (Southeast Asia Infrastructure, Jul. 21, 2021).

¹³The final locations of all stations on a specific line (or a subset of it) are announced all at once, instead of providing multiple options for each station's location. For example, the Land Transport Authority (LTA) recently announced the addition of seven new mass rapid transit (MRT) stations to the Thomson-East Coast Line (TEL). Source: <https://www.infrappworld.com/news/seven-new-stations-are-to-be-added-to-mass-rapid-transit-line-project-in-singapore>.

III.B. Data

To observe house trading activity of civil servants and other individuals, we collect data from three distinct sources. We then integrate these datasets with property transaction records using shared personal identifiers. The central data contribution is the ability to link universe-level private-property transactions to civil-service identities and, separately, to relatives/household members.

The first source is the Singapore Government Directory Interactive website, where we collect a list of civil servants. This list includes details such as names, designations (posts), and affiliated agencies/ministries. We collected panel data from six cohorts between 2007 and 2022. These cohorts are from the years 2007, 2009, 2011, 2013, 2015, and 2022. This approach allows us to gather more accurate information about the identity of house buyers and reduce misclassification.¹⁴ The second source of data is the public register of the Council for Estate Agencies (CEA), which provides a comprehensive list of licensed real estate salespersons (or agents).¹⁵ The CEA data includes information such as the salesperson's name, affiliated estate agency/firm, and registration number. Finally, we obtain the third database from a proprietary source, which contains demographic details and residence addresses for a large sample of Singaporean residents as of 2013.

We also collect property transaction data from the legal/caveat filings of private property transactions with the Registry of Land Titles between January 1995 and December 2019. More than 80% of Singaporeans live in the public housing built and allocated by the Housing and Development Board (HDB).¹⁶ The public housing segment is highly regulated and the new HDB flats are sold at subsidized prices, subject to meeting income ceiling eligibility criteria.¹⁷ In our empirical analysis, we primarily focus on the private market segment for a couple of reasons. Firstly, an unregulated housing market may provide easier access for individuals aiming to leverage private information to profit from housing transactions. Secondly, and relatedly, since private housing prices are predominantly influenced by free market forces, we anticipate them to reflect private information once it becomes public, thus potentially motivating informed acquisition prior to public project announcements.

¹⁴In contrast, exploiting a single cross-section of public officer profiles may increase both Type 1 and Type 2 errors. This means some individuals may be misclassified as civil servants (or not) at the time of transactions when they may have already left or just joined the public service sector.

¹⁵Council for Estate Agencies (CEA) is a statutory board established on 22 October 2010 under the Estate Agents Act, and is empowered with the responsibilities to enhance professionalism in the real estate agency industry, through licensing of real estate agents and also enforce fair practices with an enhanced regulatory framework.

¹⁶HDB is the public housing authority of Singapore established on 1 February 1960 with the main responsibilities of building affordable and quality public housing for Singaporean citizens.

¹⁷HDB flats are only sold to Singaporean households with an income ceiling of less than \$12,000 per month. The previous household income ceiling was set at \$10,000 per month, before it was revised upward with effect from September 2015.

The transaction data contain detailed information on housing attributes, such as property name, property type (apartment or condominium), floor, unit size, sale type (new sale, subsale and resale¹⁸); and sale details, such as date of sales, names of buyers and sellers, and sale price. Using the names, we can match them with datasets of personal identities (i.e., civil servants and the population database) to obtain information on the profession, demographics, and home address details for each buyer and seller.

A central challenge in detecting news-based insider trading is identifying the asset in which informed traders are most likely to profit before the announcement.¹⁹ This process is more straightforward in our setting. The value of MRT proximity is likely to decline with geographic distance from the station. Our empirical setting relies on the houses that have proximity benefits to the new stations. We use GIS tools to geocode MRT stations and properties and define circular buffers of radius X kilometers around each proposed station.

We classify houses as “target” (or “nearby”) properties if the houses sold fall within the X km boundary. Figure 2 illustrates the approach using the geographic distribution of residential properties surrounding Kembangan Station as an example. In our main analysis, we use a 1 km radius distance measured from the center of the proposed MRT stations as a cutoff. For robustness, we also consider other cutoffs to separate houses that are more likely to experience price increases due to their proximity to the proposed MRT stations, from those that do not have such advantages (those located outside a 1.5 or 2 km radius).

III.C. Matched Sample

While public equity investments are primarily motivated by future dividend and price expectations, as well as risk diversification benefits, the decision to purchase a house can be influenced by various factors, given the dual nature of housing as both a consumption and an investment good. First, household preferences for consumption goods, career considerations, and life cycle effects, such as decisions related to childbirth, may vary significantly. Second, unanticipated changes in wealth and income shocks can also influence individuals’ decisions to buy houses. For example, if civil servants (the treatment group) have higher average incomes compared to individuals in the control group, they may have a greater likelihood of affording more (expensive) houses. Hence, the disparities in cross-sectional characteristics between civil servants and potential control groups might contribute to the divergence in the propensity to purchase houses.

¹⁸The sale data are divided into three types according to the transaction time. New sales cover all first-hand units, usually uncompleted, that are bought directly from developers. Sub-sales are those houses that are resold before completion by individual buyers in the open market. Resale units are completed units that are sold by owners / investors after some years of occupation or holdings.

¹⁹For example, [Bernile et al. \(2016\)](#) use E-mini S&P 500 futures to examine the abnormal order balance prior to the Federal Open Market Committee’s scheduled announcements.

Our primary focus is to examine the discrepancy in the target property purchases near MRT stations, rather than overall housing consumption, which makes it less susceptible to the aforementioned concerns. Specifically, the identifying variation comes from changes in target-property purchase intensity around announcement dates, comparing civil servants to observably similar controls, rather than from cross-sectional differences in overall housing demand. Nevertheless, to further account for the potential impact of confounding factors, we employ the propensity score matching (PSM) method. Specifically, we define the treatment group as civil servants and construct a matched control group composed of non-civil servant individuals who share similar observable characteristics.

Given the sizable sample of “uninformed” buyers, we adopt a stringent matching strategy that carefully considers demographic attributes and housing consumption behaviors, aiming to closely resemble the treatment group of buyers (civil servants). We implement a two-stage matching process. In the initial stage, we employ the propensity score matching method, utilizing a comprehensive set of factors, including age, gender, marital status, and ethnic group, to account for the effects of life cycle housing consumption for the two matched samples of buyers.

In the subsequent step, we select a comparable non-civil servant (uninformed buyer) for each civil servant, ensuring that the selected non-civil servant (1) has the closest propensity score within a caliper width of 0.1 and (2) resides in the same building, as determined by sharing the same 6-digit postal code of their residential address. If a comparable control falls outside the caliper width of 0.1, we expand the sample to include non-civil servants living in the same community. Throughout the matching process, we ensure that each control sample is matched only once to a treatment sample by performing the matching without replacement.

III.D. Purchase Intensity of Target Properties

One advantage of studying informed trading in the real estate markets is that transactions are uniquely identifiable, as they are recorded based on buyer names in the transaction records. An additional advantage of our framework lies in its ability to detect such behavior among civil servants across the entire spectrum, irrespective of whether it has been officially unearthed. This sets us apart from studies identifying misconduct using solely formally discovered cases.

We concentrate on the annual intensity of transactions per person for target properties over a specific time period similar to [Cheng et al. \(2014\)](#). Transaction intensity is defined as the number of target property transactions in group i in year t , divided by the number of individuals in that group who could make such a purchase in year t . We pool all MRT events and construct a balanced panel spanning from four years before to three years after the announcement year (-4 to $+3$). The purchase intensity for group or an individual i in year t is defined as follows:

$$(5) \quad \text{Purchase Intensity}_{i,t} = \frac{\# \text{ target property transactions in group } i, t}{\# \text{ people in group } i, t}$$

IV. Informed Purchases Before Public Project Announcements

IV.A. Purchase Intensity Around the Announcements

We begin our analysis by documenting house purchases made by public officers in the vicinity of newly announced MRT stations. To formally assess differences while accounting for observable demographics and location-specific characteristics between civil servants and a matched sample of other individuals, we construct a balanced person-year panel. This panel assigns the number of house purchases to individual-event-year observations. We then estimate regressions to examine how the intensity of target property purchases varies across buyer groups and across time relative to the project announcement year.

Our baseline empirical design comparing the purchase intensity of the civil servant group with the matched uninformed buyers surrounding the announcement date can be formally represented as:

$$(6) \quad Y_{i,t} = \alpha \times Gov_i + \sum_{k=-4}^3 \beta_k \times Gov_i \times D(t = k) + \theta X_i + f_t + g_j + \varepsilon_{i,t}$$

where $Y_{i,t}$ is the individual-event-year purchase intensity defined in (5), multiplied by 100. It measures the number of target-property purchases by individual i in event year t per 100 person-event-years. The time index t is measured in event time, where $t = 0$ corresponds to the year of the MRT station announcement, negative values indicate years before the announcement, and positive values indicate years after. As very few people make more than one purchase within a year, this purchase intensity is close to an annual purchase probability. Since the outcome is multiplied by 100, the coefficients can be interpreted approximately as percentage-point changes in annual purchase intensity.

Gov_i is a dummy variable indicating whether the buyer is a civil servant or not. $D(t = k)$ is a time dummy variable that has a value of one, if the time is within k th year before the announcement, or $[k, k + 1]$. X_i is a vector of the buyer covariates including age, ethnicity, marital status, and a high/low income proxy based on the buyer's residential property type; f_t and g_j are vectors of the event time and postal code fixed effects. The main variables of interest are the estimates of the interactions between Gov_i and the time dummies $D(t = k)$, β_k , which capture changes over time in the difference in target property purchase intensity between civil servants and the control group relative to the announcement year.

Table 1 presents descriptive evidence together with the baseline event-study estimates. Panel (a) shows that matched civil servants and controls are broadly similar in age, gender, marital status, ethnicity, and approximated monthly income.²⁰ Panel (b) summarizes raw annual purchase intensity

²⁰Because individual income is not directly observed, we approximate it using a proprietary consumer financial transaction dataset from Agarwal and Qian (2014): we first compute mean monthly income by the first four digits of

of targeted properties before and after MRT station announcements within a 1 km radius of the new station. Civil servants purchase targeted properties more intensively than matched controls before the announcement, with a gap of 0.53 percentage points; the corresponding post-announcement gap is smaller, at 0.33 percentage points.

Panel (c) reports the main results using the baseline empirical design described above. Columns (1)–(2) show the results for property acquisitions within a 1 km radius of the new station. Civil servants exhibit a notable increase in purchasing target properties one to two years prior to the announcement, by 0.67 and 0.53 percentage points, respectively, compared to the matched control group. The effect is both statistically and economically significant and it corresponds to a relative increase of more than 60 percent relative to the unconditional average. The differences in purchase intensities between the civil servants and control samples are smaller and not statistically significant in other window periods, including the three years prior to the announcement and the post-announcement periods. The estimated coefficients remain largely unchanged, and the results are robust when we include buyers’ demographic characteristics and postal code fixed effects. We also observe a relative increase in the purchase intensity by civil servants for alternative choices of distance cutoffs—within 1.5 km and 2 km—compared to the baseline cutoff of 1 km (Columns (3)–(6) of Table 1).

Figure 3 visualizes these results by plotting the estimated change in target-property purchase intensity of civil servants relative to the control group around the announcement, using alternative distance cutoffs. The vertical axis shows the estimated differences in percentage points, while the horizontal axis indicates years relative to the announcement. Because the dependent variable is multiplied by 100, a coefficient of 0.5 corresponds approximately to a 0.5 percentage-point increase in annual target-property purchase intensity.

We observe a significant relative increase in property purchases near the forthcoming MRT station by civil servants in the one to two years leading up to the announcement, with the peak estimate corresponding to more than a 60% increase relative to the unconditional average. Because we do not observe exactly when each official learns station-location information, this effect should be interpreted as a local, reduced-form estimate of potential access to nonpublic information, and likely as a lower bound for officials who actually had access at the relevant time. Even this local estimate is economically meaningful: a back-of-the-envelope calculation yields approximately SGD 138 million in gross target-property acquisition value associated with the excess pre-announcement purchases identified by our design.²¹ Following the public announcement, the relative differences dissipate,

residential postal codes, and then assign this measure to individuals based on their residential postal codes. See also Section V.C.2.

²¹To gauge the aggregate magnitude, we use the balanced regression sample, which contains 9,347 civil-servant person-event observations in each event year. Applying the event-study coefficients at $t = -2$ and $t = -1$ —0.535

with civil servants and the control group exhibiting similar trading behavior. These patterns suggest that civil servants engaged in pre-emptive property purchases in anticipation of the new railway station announcement.

IV.B. Robustness Checks

Permutation Exercise. We perform two falsification tests to further examine the robustness of the above results. First, we conduct a permutation exercise by randomly assigning the dummy for civil servant to assess the validity of our political connection measure and exclude the possibility that we establish a spurious relationship. As previously shown, the evidence of insider trading becomes apparent one or two years prior to the announcement. We recode the event year dummy by combining these two event years into a new dummy and re-estimate equation (6) based on this randomization to obtain the $\beta_{-1,-2}$ coefficient estimate for $D(t = -2 \text{ or } -1)$. We perform 1,000 random permutations and generate 1,000 corresponding bootstrap simulations. The p -value is then computed by locating the absolute value of the true estimate within the distribution of simulated estimates under the null hypothesis. Second, we conduct a complementary falsification test by jointly randomizing MRT station locations and announcement timing. For each of 1,000 simulations, we redraw MRT station locations within Singapore, assign placebo announcement months, rebuild the event panel, and re-estimate equation (6) to obtain the corresponding $\beta_{-1,-2}$ coefficient estimate.

In Figure 4, we plot the distribution of the bootstrapped estimate, $\beta_{-1,-2}$, with the fully saturated specification (as in Columns (2), (4), and (6) of Table 1) and three different distance cutoffs. Panel (a) reports the permutation exercise that randomizes civil servant status, while Panel (b) reports the exercise that jointly randomizes MRT station locations and announcement timing. The coefficients peak at zero and exhibit a symmetrical distribution on both sides. In the figure, we also plot the coefficient of the true estimate using the vertical solid line. We find that the real case coefficients across specifications have a p -value below 0.005 in the bootstrapped distribution, which is highly significant. In other words, our bootstrap analysis further confirms that the evidence of informed trading by civil servants is unlikely to hold by chance.

Placebo Tests using HDB Houses. Next, we conduct an additional placebo test using transactions in Singapore’s subsidized public-housing market. While purchasing an HDB resale flat generally takes a shorter waiting time compared to buying new HDB flats, there are still specific eligibility criteria and schemes that must be fulfilled for an individual to qualify for an HDB resale flat. These restrictions make HDB flats a less favorable vehicle for informed trading in our context. To be eligible for an HDB resale flat, an individual must (1) not own any other residential properties or any

and 0.674 purchases per 100 person-event-years—to these observations implies approximately 50 and 63 additional target-property purchase incidences, respectively, or about 113 across the two years preceding public announcements. Valuing these incidences at the mean purchase price reported for informed transactions in Section V.A.2 (SGD 1,224,150) yields an implied gross acquisition value of approximately SGD 138 million.

houses, buildings, or land (either locally or overseas) or have disposed of such properties within 6 months before the resale flat application; (2) be within the Ethnic Integration Policy and Singapore Permanent Resident quota of the block/neighborhood they wish to purchase a flat in.²²

If the increase in nearby private-property purchases documented in Table 1 reflects early access to nonpublic station-location information, we would expect much weaker effects in this more regulated market, where eligibility and resale constraints limit the scope for monetizing such information. To examine this, we replicate the analysis depicted in Figure 3 by comparing the purchase intensity of nearby public-housing units among civil servants and their matched controls around MRT station announcements. Consistent with this interpretation, Figure 5 shows no evidence of increased target-property purchase activity by civil servants relative to matched controls across alternative distance cutoffs.

This placebo exercise therefore reinforces the interpretation that the private-housing results reflect informed trading rather than a general increase in civil servants' demand for housing near future MRT stations.

IV.C. Heterogeneity in Insider Purchase – Rank, Agency, and Demographics

We next examine the distribution of misconduct across the civil-service hierarchy to address the first question posed by the model. Specifically, we analyze how the likelihood of purchasing target properties before public announcements varies across high-, mid-, and low-ranking officials. We then examine whether misconduct is more prevalent among officials employed in government agencies involved in rail-line planning, where access to nonpublic information is likely to be greater. Finally, we investigate whether the propensity to engage in informed purchases varies systematically with the demographic characteristics of civil servants.

IV.C.1. Variation Across Rank and Agency

Position Rank — We first assess how the civil servants' rank—high, mid and low—affects the likelihood of purchasing the target properties prior to the announcement. Since rank is not directly available in the data, we use an LLM-assisted classification protocol to map civil-service titles and agency information into high-, middle-, and low-rank categories. Appendix B describes the classification rules, prompt structure, and validation procedure. High-rank civil servants include senior officials whose positions are director-level or above, including political appointees, commissioners, commanders, and heads of agencies or divisions. Mid-rank civil servants are those in roles such as deputy director, assistant director, or deputy head. Low-rank civil servants include positions such as executive, engineer, or analyst. We then decompose the Gov_i dummy into

²²Source: <https://www.99.co/singapore/insider/buying-hdb-resale-flat-eligibility/>.

high, mid, and low rank indicators²³ and repeat the analysis from Table 1, combining event years, $D(t = -2 \text{ or } -1)$, into a single pre-announcement indicator.

The results are presented in Column (1) of Table 2 and in panel (a) of Figure 6. Relative to matched uninformed buyers, mid-rank civil servants exhibit a statistically significant increase in target-property purchases during the one to two years before the announcement. Mid-ranking civil servants show a meaningful and statistically significant increase in target property purchases prior to the announcement, while the effect for low-ranking officials is negligible and not statistically significant. For high-ranking officials, the estimated effect is larger in magnitude than for other groups, but it is not statistically significant.²⁴ These findings indicate that the informed-purchase pattern documented in Table 1 is primarily driven by mid-level civil servants.

As predicted by the model, this pattern supports the middle-concentration regime. Low and high ranks are effectively cornered—either because illicit returns are too small or because expected sanctions and reputational costs are too high—so only the middle rank remains responsive to enforcement. Indeed, while all levels of civil servants may have access to such information in advance, mid-level civil servants may be more vulnerable to the economic temptations of insider trading. They often have the financial means to act on the information but are not subject to the same level of public scrutiny as Ministers or senior management in government ministries and agencies. Furthermore, given the significantly higher wealth of high-ranking civil servants, the advantages of insider purchases may be relatively diminished for them compared to their mid-level counterparts.

Proximity to the Source — We also investigate the impact of proximity to the information source on the occurrence of informed trading by civil servants. To accomplish this, we focus on the purchasing behavior of civil servants employed in government agencies involved in the planning of rail line systems. Although this analysis helps establish a link between access to information and the purchase of target properties, the theoretical prediction remains ambiguous *ex ante*. On one hand, civil servants working in these agencies may have early access to the information and therefore be more likely to engage in informed purchases. On the other hand, there may be specific rules or regulations prohibiting such activity, or the associated reputation costs may be high, resulting in a lower likelihood of observing such behavior.

To assess this, we decompose the Gov_i dummy into two indicators: one for employees working in planning departments (Gov-Planning) and another for employees working in other government

²³The high, mid and low rank civil servants consist of 19%, 39%, and 42% of the treatment sample.

²⁴We also conduct a formal test for heterogeneous impacts across public officers of different ranks. In Appendix D1, Panel (a), we find that the difference in coefficient estimates between mid- and low-rank officers is statistically significant at the 1% level for both the full sample and the subsample of earlier events. In Panel (b), we do not reject the null hypothesis that the difference between mid- and high-rank officers is zero. However, this result should be interpreted with caution, as the small sample size for high-rank officers may result in low statistical power and imprecise estimates.

departments (Gov-Others). The results in Column (2) of Table 2 and Panel (b) of Figure 6 indicate that civil servants from planning agencies exhibit substantially higher target property purchase intensity compared to matched control groups prior to the announcement. In contrast, the purchase activity of civil servants in other agencies is economically smaller and statistically insignificant.²⁵

We further examine the interaction between rank and agency by estimating informed-purchase patterns separately for high-, middle-, and low-rank officials in planning and non-planning agencies. The estimates are presented in Panel (c) of Figure 6. Consistent with earlier findings, we observe that the significant increase in purchases of target properties by civil servants one to two years before the announcement is primarily driven by mid-level civil servants in planning departments. This additional evidence reinforces the interpretation that such purchases likely reflect informed trading.

IV.C.2. Demographics

Demographic characteristics can influence the likelihood of informed purchase activity. For example, trading behavior may differ across households at various stages of the life cycle: older male civil servants, who typically have greater savings and accumulated wealth, may be more inclined to engage in informed purchases. To examine this heterogeneity, we categorize buyers into high and low age groups, as well as by gender and ethnic group, and estimate the purchase intensity model using interactions of these group indicators with the civil servant dummy and time dummies.

The results, presented in Columns (3)-(5) of Table 2, support the notion that informed purchase intensity is higher among older civil servants, potentially reflecting higher income levels. Male civil servants also exhibit a higher level of informed purchases compared to females. However, we find little evidence that civil servants from specific ethnic groups—Chinese, Malay, or Indian—are significantly more likely to engage in informed purchases than others.

IV.D. Beyond Civil Servants: Information Advantage from Market Expertise?

Our analysis thus far provides evidence consistent with informed purchasing by civil servants: they increase their purchases of homes near planned MRT stations before those station locations are publicly announced. An alternative explanation is that the difference between civil servants and other buyers reflects superior processing of public information rather than access to nonpublic information.

To investigate this explanation further, we conduct two additional tests. Specifically, we turn our attention to two other groups of buyers in the housing market, real estate agents and corporate board members. These groups may be regarded as possessing financial expertise and the potential ability

²⁵We formally test whether the difference in coefficient estimates between officers at planning and non-planning agencies is significantly positive in Appendix D1, Panel (c). Consistent with the patterns observed in Figure 6, we reject the null hypothesis that this difference is zero at conventional significance levels for both the full sample and the subsample of earlier events.

to infer information about forthcoming public projects. We examine whether elevated purchases of homes near planned station locations are unique to civil servants or also appear among these other sophisticated buyers. Notably, this also serves as a joint test: if there is leakage of information from civil servants to these groups, we may observe limited differences in trading behavior between them, even if the original source of the information advantage lies with civil servants. Appendix D2 reports descriptive statistics for the matched samples of real estate agents and corporate board members, including observables and annual purchase intensity before and after MRT station announcements.

IV.D.1. Real Estate Agents

Real estate agents play an important intermediary role in housing markets. They use their local market knowledge to bridge the information gap between potential buyers and sellers. In this section we formally test whether real estate agents also purchase nearby properties at a higher rate. Specifically, we follow the same procedure in Section III.C to create a matched sample of real estate agents and general buyers (or uninformed buyers) and estimate the specification (6). The results, shown by the blue circles in Figure 7, present the estimated dynamic effects in an event-study plot where we include interaction term of $Agent_i$ dummy with time dummies $D(t = k)$ along with demographic controls and spatial and event year fixed effects.

Unlike the pattern in Figure 3, real-estate agents do not exhibit higher purchase intensity near planned stations than matched control buyers during the one to two years before the announcement. As can be seen, the corresponding point estimates are statistically insignificant. Overall, we find no evidence that real-estate agents display the elevated pre-announcement purchases of target properties observed among civil servants.

IV.D.2. Corporate Board Members

We turn to another group of financially sophisticated buyers in the real estate market: corporate board members. We follow the same procedure as in Section III.C to construct a matched sample of uninformed buyers (controls) for the connected board members. Then we construct a balanced person-year panel, where we assign the number of house purchases to an individual-year observation for the individual i in year t and run regressions to examine variations in the intensity of target purchases by the two buyer groups in different years prior to and after the event date. The results are shown by orange triangles in Figure 7. The coefficients on the interaction terms for all periods, $Board_i \times D(t = k)$, are statistically insignificant, which indicates that connected board members do not have a higher marginal intensity to purchase target properties compared to (uninformed) buyers.

Among board members, access to policy-relevant information may vary with their network breadth and firm characteristics. For example, the effect may be stronger for large firms which are more likely to attract important politicians to serve as the board member. The average null effect might mask significant heterogeneity, as board members' ability to access non-public government information can vary due to several factors. For instance, serving on multiple boards could potentially widen

one's access to both public and non-public information. We first test whether board members serving on multiple boards exhibit higher pre-announcement purchase intensity, consistent with broader access to information. We decompose the $Board_i$ dummy into $Board-Multi_i$ and $Board-Single_i$ dummy in the interaction term with $D(t = k)$ and re-estimate the specification. Likewise, board members of large firms may have better information access due to their political connections. We also classify firms as large or small based on whether their market value is above the sample median and repeat the analysis.

The results shown in Appendix D3 indicate that relative to matched controls, the insider purchase intensity of corporate board members is not significantly higher than that of matched controls prior to the announcement, also when we consider the type of board and the size of the firm (large vs small).

Overall, this evidence supports the view that increased purchase intensity of nearby homes by civil servants prior to project announcements reflects their access to non-public information rather than their ability to strategize and time the market based on public information. These results also indicate that there is limited, if any, information leakage from civil servants to other financially sophisticated buyers, such as board members or real estate agents.

V. Economic Gains from Insider Trading and Its Broader Consequences

V.A. Economic Gains from Insider Trading

Real estate can be purchased with the expectation of attaining a return (an investment good), with the intention of using it (a consumption good), or both. This duality complicates the process of quantifying the economic gains associated with informed trading within our context.²⁶ In this section, we attempt to gauge the extent to which the informed purchases are driven by the desire for monetary gain by examining their purchase prices, realized returns, and holding periods.

V.A.1. Costs of Informed Purchases

Unlike the stock market, housing market transactions involve direct negotiations between buyers and sellers. To the extent that civil servants' informed purchases are associated with higher levels of information asymmetry, the sellers may charge civil servants significant premiums compared to other buyers. These price premiums can potentially diminish the information advantage of civil

²⁶Living in close proximity to a Mass Rapid Transit (MRT) station offers several consumption-related benefits. First, it substantially eases commuting. In 2010 and 2020, 54.6% and 57.7% of employed residents, respectively, used buses, MRT, or Light Rail Transit (LRT) as their usual mode of commuting to work (Singapore Department of Statistics, 2021). In addition, MRT proximity often improves access to amenities and services, as many MRT stations are adjacent to or integrated with shopping malls and mixed-use developments, providing convenient access to retail, dining, and other essential services.

servants. To investigate this, we estimate hedonic housing-price models and examine whether civil servants pay higher prices for target properties than non-civil servants around announcement dates.

We consider a broader distance cutoff (within 2 km from the proposed MRT stations) to obtain a reasonable level of power for the statistical test. We control for a set of variables, including property size, a first-floor indicator, an indicator for condominium units, and an indicator for whether the buyer used central provident fund (CPF) to purchase the property. Private houses with freehold tenure command positive premiums relative to leasehold houses (Giglio et al. 2015) and thus we include the Freehold dummy control for the underlying property.²⁷ The year and the spatial (6-digit postal code) fixed effects are also included in the specifications. Again, we focus on the interactive variable, $Gov_i \times D(t = k)$, which captures the price differences in the transactions by civil servants and non-civil servants over the target properties around the announcement.

Table 3 presents the results of the hedonic price model that compares the purchase prices paid by civil servants with that by the other two control groups for target-property purchases. The results do not support the hypothesis that sellers charge civil servants significant premiums relative to other buyers. As can be observed, civil servants purchase the target property at a similar price relative to the matched uninformed buyer and real estate agent. The only significant estimate is the unadjusted total-price interaction at $t = -2/-1$ (8.5 percent), which falls to 2.4 percent and becomes insignificant after controlling for price-relevant observables such as size. The corresponding estimates for log price per square meter are also small and statistically insignificant. Collectively, these results reject the hypothesis that civil servants pay significant premiums on their “informed” purchases relative to other buyers.

V.A.2. Realized Returns from Informed Trading

We next examine the performance of informed purchases made by public officers. Specifically, we focus on consecutive transactions—that is, instances where informed purchases are followed by a corresponding sale—allowing us to calculate the realized returns earned between the purchase and sale dates. For that purpose we define the unlevered annualized return for the property i bought in year t_0 at the price of P_{i,t_0} and sold in year t_1 for the price of P_{i,t_1} as follows:²⁸

$$(7) \quad r_{i,t_0 \rightarrow t_1}^{unlev} = \left(\frac{P_{i,t_1}}{P_{i,t_0}} \right)^{\frac{1}{t_1 - t_0}} - 1$$

²⁷Following He et al. (2019), which shows that houses with a lucky unit number ending with “8” also command a premium, especially for Chinese buyers, relative to houses with an unlucky unit number ending with “4”, we include units and floor levels with numbers “4” and “8” as additional control variables in the hedonic price models.

²⁸This is the gross return, excluding intermediate factors such as home renovation costs, for which we do not have data. It also does not account for other costs and benefits of homeownership, including ongoing ownership expenses or savings from not having to pay rent or from not living in an alternative property. However, this measure remains informative, particularly if these omitted factors do not differ systematically relative to the matched sample of other buyers.

In practice, returns from real estate investment are typically leveraged, as a significant share of homeowners finance their home purchases through mortgages. This leverage is generally sustained over long periods due to the nature of long-term mortgages, which feature fixed amortization schedules that front-load interest payments. However, the dataset does not include detailed mortgage information—such as term length, interest rate, or down payment—needed to precisely calculate leveraged returns for individual homeowners. To address this, we follow the approach of [Goldsmith-Pinkham and Shue \(2023\)](#) and adopt a simple benchmark scenario: we assume homeowners use the most common mortgage type in the dataset—a 30-year fixed-rate loan with an initial loan-to-value (LTV) ratio of 80% and prevailing market interest rates at the time of purchase—without refinancing.

We use a similar formula as in equation (7) to calculate the levered annualized return, where $Equity_{i,t_1}$ is the total cash out at sales time t_1 and $Equity_{i,t_0}$ is the net present value of equity at purchase time t_0 :

$$(8) \quad r_{i,t_0 \rightarrow t_1}^{lev} = \left(\frac{Equity_{i,t_1}}{Equity_{i,t_0}} \right)^{\frac{1}{t_1 - t_0}} - 1$$

To operationalize this formula, we begin by using the purchase value and assuming an 80% loan-to-value (LTV) ratio to calculate the initial down payment and mortgage amount for the property acquisition. Next, we determine the interest rate corresponding to the transaction’s year and month, based on the long-term mortgage rates published by the Monetary Authority of Singapore. This interest rate is then used to estimate the reduction in principal over time—at each monthly interval—assuming no refinancing. From this, we can identify the proportion of the outstanding mortgage principal at the point of sale in period t_1 , M_{i,t_1} . Third, $Equity_{i,t_1}$ is calculated as the maximum of $(P_{i,t_1} - M_{i,t_1})$ and zero. The equity at purchase $Equity_{i,t_0}$ is the sum of down payment plus discounted value of all principal repayments over time of ownership.

Panel (a) of Figure 8 shows raw unlevered and levered annualized returns at various percentiles of the return distribution for informed purchases by public officers. The panel reveals that these transactions on average generate positive returns in all parts of the return distribution except the left tail. For instance, the mean and median of the unlevered annual return are 12% and 5%, respectively, while the mean realized levered return is 57.9%.

Next, we examine whether the observed positive returns are attributable to public officers’ information advantage or simply a reflection of the broader housing market in Singapore. To explore this, we construct benchmark returns using two approaches. First, in Panel (b) of Figure 8, we adjust for market trends by calculating the value-weighted return of all repeated sales that occurred in the same purchase and sale years (i.e., purchased in year “ t_0 ” and sold in year “ t_1 ”). Second, in Panel (c) of Figure 8, we adjust for similar purchases in different locations by using the value-weighted return of all repeated sales that took place in the same purchase and sale years but outside the target

regions (i.e., purchased in year “ t_0 ” and sold in year “ t_1 ” but located outside the MRT station’s neighborhood).

Panels (b) and (c) of Figure 8 display the distribution of these adjusted annualized returns by showing the difference between the returns on target properties purchased by civil servants and the market returns, as calculated in the manner described above. The figure confirms that informed purchases outperform benchmark housing returns across most of the return distribution, with the largest differences observed in the right tail. For instance, informed buyers earn annualized returns that are, on average, 2.4 percentage points higher on an unlevered basis and 16.3 percentage points higher on a levered basis relative to the market benchmark. When benchmarked against non-target properties—those located farther from railway stations—the performance gap widens to 6.8 and 36.2 percentage points for unlevered and levered returns, respectively.

We also note that the gains from informed trading are substantial. As depicted in panel (a) of Figure 8, the average annualized return is about 12%. Given the average housing value of these informed transactions is 1,224,150 SGD, this translates to additional annual earnings of 146,898 SGD, which is approximately an average annual income of these civil servants (136,596 SGD²⁹). Even when accounting for market trends or similar trades in control regions (as shown in panel (b) and (c)), the extra return is economically meaningful and equivalent to 21.5% ($2.4\% \times 1,224,150 / 136,596$) and 60.9% ($6.8\% \times 1,224,150 / 136,596$) of their average annual income, respectively.

V.A.3. Holding Period

We provide additional evidence regarding the investment incentive of informed purchases by examining their holding periods. Holding consumption motives fixed, investment-driven purchases are more likely to be sold sooner. We track subsequent sales, when observed, for properties purchased by civil servants around announcement dates. We then calculate the holding periods (in months) for each informed purchase in our sample. As in the previous section, we also compute the benchmark holding period for similar purchases in different locations. Panel (a) of Figure 9 shows both mean and percentiles of holding periods for both types of transactions. We observe that civil servant buyers tend to keep properties for slightly shorter periods if the properties were purchased one or two years before the announcement.

To formally evaluate whether informed buyers sell their properties faster, we also conduct a survival analysis with Kaplan-Meier estimators to model the time-to-sale over a 10-year horizon post-purchase and plot the survival curve in Panel (b) of Figure 9. Specifically, we track the probability of houses “surviving” unsold over time between informed purchases (by public officers) and uninformed purchases (occurring at the same time in non-target regions). We find that the

²⁹As shown in Table 4, these civil servants earn an average monthly income of 11,383 SGD translating into annual income of 136,596 SGD.

survival probability for informed purchases consistently falls below that of their counterparts at every observed interval, especially during the first six years since purchase. The results suggest that informed buyers likely liquidate assets faster due to investment incentives.

In conclusion, this evidence suggests that target property purchases by civil servants may be driven by economic motives, as they are associated with higher returns and shorter holding periods. However, these findings should be interpreted with caution due to the relatively small number of ex-post sales observations in our sample.

V.B. Spillover Channel: Relatives' Informed Purchases and Gains

Insider trading often involves information leakage and the participation of related parties. For example, the 2025 resignation of a Fed governor in the US has been associated with allegations of insider trading by their relatives. In our setting, informed civil servants may share information with close family members or friends. These connected individuals may either substitute for the official by transacting in their own names or coordinate purchases alongside the official. Motivated by these possibilities, this subsection examines whether relatives engage in informed purchases around station announcements and assesses the associated economic gains.

Because we do not observe family ties directly, we identify connected individuals using administrative co-registration at the same property as the civil servant; that is, individuals listed as co-occupants or co-registrants at the same residential address. We refer to these individuals as “relatives” for expositional convenience. This proxy is useful in Singapore because residential identifiers are highly granular, allowing us to identify individuals linked to the same property rather than merely to the same neighborhood. Co-registered individuals are therefore plausible recipients of privileged information and a natural channel through which such information may affect trading by connected individuals.

To examine this channel, we match the relatives sample to a control group and estimate specification (6) for the pre-2011 period. Panel (a) of Figure 10 presents the dynamic estimates and shows a pronounced pre-announcement increase in relatives' purchase intensity compared with matched controls, consistent with the use of non-public information. We refer to this mechanism as the spillover (indirect) channel of informed trading. This spillover channel is economically meaningful: a back-of-the-envelope calculation yields approximately SGD 132 million in gross target-property acquisition value associated with the excess pre-announcement purchases by relatives identified by our design.³⁰

³⁰To gauge the aggregate magnitude, we use the implied balanced regression sample, which contains 13,834 relative person-event observations in each event year. Applying the event-study coefficients at $t = -3$, $t = -2$, and $t = -1$ —0.311, 0.224, and 0.246 purchases per 100 person-event-years—to these observations implies approximately 43, 31, and 34 additional target-property purchase incidences, respectively, or about 108 across the three years preceding public

We next evaluate the economic payoffs and holding horizons of relatives' trades. Panel (b) of Figure 10 plots unlevered and levered annualized returns across percentiles for relatives' informed purchases. Returns are positive and economically sizable across most of the distribution. Appendix D4 corroborates this by replicating the analysis underlying Figure 8 and showing that relatives' realized returns are, on average, significantly above both market benchmarks and non-target controls, for both unlevered and leveraged measures. Kaplan-Meier survival estimates in panel (b) of Appendix D5 show that properties acquired by relatives are more likely to be sold earlier—over horizons from one to eight years after purchase—than comparable control transactions, consistent with investment-driven behavior and timely exit.

Relatives, unlike civil servants, also display substantial post-announcement buying. Coefficients for event times from $t = 0$ to $t = +3$ in Figure 10, panel (a) are generally positive and statistically significant. Two forces can explain this pattern. First, relatives and close friends often prefer to live near one another, which raises baseline demand in places where civil servants buy. Second, relatives who receive timely information may process and act on the station announcement faster than other buyers. Acting before the information is fully capitalized into prices may allow them either to earn higher returns or to purchase high-amenity homes near the new station at relatively favorable prices. Although we do not report full results, two additional results align with both channels: (i) relatives' post-announcement purchases continue to earn returns significantly above the market and matched controls; and (ii) unlike pre-announcement purchases, Kaplan-Meier hazard estimates indicate longer holding periods over one- to five-year horizons relative to controls. The latter suggests many such buyers are amenity-motivated rather than just pursuing short-term speculative gains.

Taken together, the evidence establishes a meaningful spillover (indirect) channel: relatives of civil servants also engage in informed purchases and earn economically sizable gains from doing so. This broadens the scope of informed trading beyond direct purchases by officials and highlights the importance of connected individuals for both measurement and enforcement design.

V.C. Broader Consequences of Insider Property Purchases

In this subsection, we discuss and analyze the broader implications of the informed trading by civil servants that we documented above, including some potential welfare consequences. The presence of such informed trading can lead to welfare changes through at least three channels.

V.C.1. Externalities to the Market

First, the insider purchases may introduce negative externalities that impact the allocative efficiency of the market. The significant information asymmetry between insiders and outsiders can lead to limited trading and reduced market liquidity, as uninformed participants may hesitate to engage in

announcements. Valuing these incidences at the mean purchase price used in the civil-servant calculation in Section V.A.2 (SGD 1,224,150) yields an implied gross acquisition value of approximately SGD 132 million.

transactions due to uncertainty (Akerlof 1970). While this issue is relevant to all forms of insider trading, it is challenging to identify such effects within the specific context of our study.

V.C.2. Redistributive Effects

Second, the purchase of properties by civil servants prior to the announcement of MRT station locations may reduce the welfare of uninformed agents, particularly sellers, resulting in a wealth transfer from uninformed sellers to informed buyers. The impact on weighted social welfare will depend on the relative marginal utility of wealth between buyers and sellers. Specifically, the effect is more pronounced when sellers have a higher marginal utility of wealth, as the loss of wealth for these individuals has a greater negative impact on their overall well-being. In contrast, the effect is less severe when sellers have a lower marginal utility of wealth, as the wealth loss would cause a smaller reduction in their welfare.

Although marginal utility is not directly observable, we can approximate it using income levels. To do this, we use a proprietary consumer financial transaction dataset from Agarwal and Qian (2014), obtained from a major bank in Singapore. The dataset includes individual income data, which we aggregate by region (using the first four digits of the postal code) to estimate average incomes for buyers and sellers in both general housing transactions and informed purchases. Table 4 presents the results: Panel (a) covers the full sample, while Panel (b) focuses on transactions involving target properties in event years $t-1$ and $t-2$. Although buyer and seller characteristics vary across samples, their income levels are generally similar. This indicates that sellers—presumably uninformed—are not significantly richer than informed buyers. As a result, the welfare effects of transfers from sellers to civil servants engaging in informed trading may be tangible, since they are not offset by redistributions from wealthier to poorer households.

V.C.3. Suboptimal Policy

Third, social welfare may decline due to suboptimal decisions regarding station placement. Insiders may acquire properties in advance and then attempt to influence the location of new stations to increase the value of their holdings. This manipulation can lead to inefficient station locations, reducing the overall effectiveness of the railway system and ultimately harming social welfare.³¹

To investigate the potential for location manipulation, we analyze whether property purchases by public officers are more likely to predict the locations of future MRT stations compared to other buyers. We control for house characteristics, transaction details, demographic factors, transaction year, and spatial fixed effects, as specified in the hedonic price model. The results, presented in Table 5, show that the coefficient on the “Gov” dummy indicates no statistically significant difference in

³¹MRT stations may be suboptimally located, deviating from their ideal placement based on urban planning efficiency or geotechnical factors such as soil composition, groundwater levels, seismic risk, and slope stability. However, due to the absence of data on the relative cost-efficiency of alternative locations, we are unable to evaluate the extent of this potential inefficiency.

the distance to the MRT station between properties purchased by civil servants and those purchased by matched uninformed buyers. Therefore, we find little evidence to support the hypothesis that civil servants influence station location decisions *ex ante*.

VI. Formal Enforcement, Social Norms, and Misconduct

We conclude our analysis by addressing the second question posed by the model: whether, after decades of successful enforcement, norms become sufficiently persistent to sustain a high-integrity equilibrium or whether continued formal enforcement remains necessary. We study this question through the 2011–2012 enforcement tightening, which provides a natural setting for examining how misconduct responds to stronger formal enforcement. Later in this section, we relate the response of misconduct to broader developments in the integrity environment using independent evidence on perceived corruption, public trust, and pro-integrity norms. Together, these pieces of evidence allow us to assess whether the deterioration in the integrity environment reflected genuine behavioral deterioration or merely noisy signals, and therefore whether continued formal enforcement remained necessary to sustain strong norms.

Our analysis focuses on the 2011–2012 enforcement tightening, centered on the implementation of the Securities and Futures (Amendment) Act 2012. The amendment introduced major reforms to Singapore’s capital market framework, including expanded insider trading provisions, enhanced investigative powers for the Monetary Authority of Singapore (MAS), stiffer penalties for market misconduct, and the establishment of a whistle-blowing framework. In the same year, revisions to the Official Secrets Act (Chapter 213, Revised Edition 2012) broadened the scope of protected information and strengthened information security practices within government agencies. These legislative changes coincided with several high-profile corruption cases—including the 2012 scandals involving senior officers from the Central Narcotics Bureau (CNB), the Singapore Civil Defence Force (SCDF), and the Corrupt Practices Investigation Bureau (CPIB)—which brought renewed public attention to integrity in the public sector (Lee 2012).

VI.A. Enforcement Tightening and Misconduct

VI.A.1. Enforcement Tightening and Informed Purchases by Civil Servants

To investigate how misconduct changes around this enforcement tightening, we compare project announcements made on or before 2011 with those afterward. This comparison allows us to examine how misconduct responded to stronger formal enforcement.

Panel (a) of Figure 11 presents event-study plots showing the estimated differences in purchase intensity of target properties by civil servants around the announcement events. The estimates indicate that the informed trading activity documented in Section IV was largely concentrated in the period before 2011, with a notable decline thereafter. These findings are consistent with the view

that regulatory and policy changes effectively curtailed the misuse of non-public information for personal gain among civil servants. Figure 6 reinforces this conclusion by showing no evidence of insider trading by civil servants after 2011, including among mid-level officials and those in planning departments, where such behavior was evident prior to 2011 (see also Appendix D1).

The evidence so far indicates that informed purchases by civil servants decline after the legislative and enforcement changes that increased the expected costs of misusing privileged information. However, officials may respond by shifting informed purchases to connected individuals, such as family members or co-registered household members. In that case, direct purchases by officials could fall even if the overall level of informed trading remains unchanged. To explore this further, we next analyze the trading patterns of family members of civil servants.

VI.A.2. Enforcement Tightening and Informed Purchases by Relatives

As discussed in Section V.B, we infer the identities of these individuals using the population dataset by focusing on the other members who registered at the same property as each civil servant. We pair this sample of relatives with a matched control group and estimate the specification (6). The dynamic effects are illustrated in an event-study plot of panel (b) in Figure 11.

Prior to 2011, the purchase intensity of target properties by relatives was elevated and borderline significant relative to the matched controls in the one to two years before the announcement, and there is a significant uptick already at $t = -3$, complementing the timing of civil servants' purchases. These pre-announcement movements are consistent with an information spillover channel in which relatives either substitute for the civil servant or coordinate purchases with them.

Relatives, unlike civil servants, also display substantial post-announcement buying. Coefficients for event times from $t = 0$ to $t = +3$ in Figure 11, panel (b), are generally positive and statistically significant. As discussed in Section V.B, these post-announcement trades reflect a mix of amenity-driven motives and residual informational advantages. They are associated with positive abnormal returns and longer holding horizons relative to controls. Taken together with the pre-announcement evidence, relatives' post-announcement behavior in the pre-2011 period further confirms that the spillover channel operates both before and after station announcements.

Turning to the post-2011 period, the effects of regulatory and enforcement tightening are stark. The post-2011 evidence shows no comparable elevation in informed purchases: for both civil servants and their relatives, the pre- and post-announcement coefficients shrink toward zero and are statistically indistinguishable from matched controls. These patterns are also inconsistent with any substitution from officials to family members under heightened scrutiny. The absence of pre-announcement elevation suggests that civil servants no longer systematically act on, or transmit, privileged information before announcements. The muted post-announcement response among relatives further suggests that any residual informational advantage either dissipated or became less

valuable once officials themselves stopped concentrating purchases near newly announced locations. Taken together, the evidence indicates that stricter enforcement substantially reduced both the direct channel of informed trading by officials and the indirect channel through relatives.

VI.B. Broader Integrity Environment around the Enforcement Tightening

The evidence thus far shows that stronger formal enforcement substantially reduced misconduct. This naturally raises a broader question. Why did policymakers choose to strengthen formal enforcement after decades of successful anti-corruption policy? The model suggests that governments may not be able to answer this question by directly observing underlying misconduct. Instead, they rely on broader signals about the integrity environment when deciding whether norms remain sufficiently persistent or whether stronger formal enforcement is warranted. We therefore examine independent survey evidence on perceived corruption, public trust, and pro-integrity norms surrounding the 2011–2012 enforcement tightening.³² These measures provide empirical counterparts to the noisy signals emphasized in the model that policymakers may have observed when deciding whether to strengthen enforcement.

Figure 12 summarizes the three sets of survey-based measures. Panel (a) reports evidence from the [Asian Barometer Survey](#) on perceived corruption. Around the 2011–2012 enforcement episode, respondents became substantially more likely to report witnessing corruption and to perceive corruption among public officials. By 2020, both measures returned close to their earlier levels. Panel (b) reports survey evidence on public-sector trust and reputation. Distrust in public institutions, including the civil service, parliament, courts, and the executive, increased around the same period before generally declining thereafter. Panel (c) presents evidence from the [World Values Survey](#) (Inglehart et al. 2022) on pro-integrity norms. Across several forms of dishonest behavior—including accepting bribes, tax evasion, fare evasion, and fraudulent benefit claims—the share of respondents reporting that such behavior is "never justifiable" declined around 2012 before recovering by 2020. Taken together, the three sets of measures point to a common pattern: the broader integrity environment appeared to deteriorate around the enforcement tightening and subsequently recovered.

Viewed in isolation, however, these survey patterns remain fundamentally ambiguous. Higher perceived corruption, weaker pro-integrity norms, and lower trust in public institutions may reflect genuine deterioration in underlying misconduct, but they may also arise from changing perceptions,

³²We use the [World Values Survey](#) (Inglehart et al. 2022) and the [Asian Barometer Survey](#). These surveys are widely used in academic and policy research, provide publicly available data for Singapore, and contain questions relevant to our setting. They are also unusually well suited to our analysis because their Singapore waves provide observations before the 2011–2012 enforcement tightening, around the episode, and after the episode, with coverage extending from the early or mid-2000s to the early 2020s. To the best of our knowledge, they are the only publicly available repeated surveys for Singapore that jointly satisfy these timing, coverage, and question-content requirements.

increased media attention, or other sources of noise. In other words, the surveys characterize the type of information available to policymakers, but they cannot establish whether the underlying state of misconduct had actually changed. This ambiguity is precisely the government’s information problem highlighted by the model.

Unlike policymakers, however, we observe a transaction-level measure of underlying misconduct. Combined with the evidence from Section VI.A, this allows us to determine whether the deterioration in the broader integrity environment was accompanied by genuine behavioral deterioration and whether the subsequent strengthening of formal enforcement responded to an actual weakening of compliance rather than noisy signals alone. The next subsection brings these pieces together through the lens of the model and returns to the broader question of whether, even after decades of successful enforcement, norms were sufficiently persistent to sustain Singapore’s high-integrity equilibrium without continued formal enforcement.

VI.C. Inferring Norm Persistence from Misconduct and Enforcement

The model distinguishes between two interpretations of the 2011–2012 enforcement tightening. A noise-driven tightening does not require genuine deterioration in underlying misconduct. Policymakers may temporarily strengthen enforcement because they observe adverse but ultimately noisy signals about the integrity environment. By contrast, a learning-based tightening reflects an informative inference that norms are less persistent than previously believed. In this case, deterioration in the broader integrity environment should coincide with actual misconduct, and stronger formal enforcement should reduce that misconduct.

Several features of the data help distinguish between these two interpretations. First, we document meaningful pre-policy misconduct by both civil servants and their relatives, indicating that the deterioration in the integrity environment was accompanied by genuine underlying behavioral deterioration rather than merely changing perceptions. Second, the response to stronger formal enforcement is broad-based across the margins where misconduct previously appeared, including mid-level officials, planning agencies, and relatives of civil servants. This pattern is consistent with effective deterrence of actual misconduct rather than an unnecessary response to noisy signals. Finally, the decline in misconduct persists well beyond the immediate reform period rather than exhibiting a short-lived dip followed by mean reversion, suggesting a durable change in behavior rather than a temporary reaction to heightened scrutiny.

The correspondence between the survey and transaction-based evidence also clarifies how the survey measures should be interpreted. Their alignment with the observed informed-purchase activity suggests that perceptions of corruption contain information about underlying integrity conditions, even if they remain noisy. Moreover, the common movement in perceived corruption, public-sector distrust, and pro-integrity attitudes is consistent with the feedback mechanism in the

model: rising misconduct weakens trust and the reputational value of integrity, which may further erode incentives for compliance. Conversely, the post-tightening decline in misconduct, together with the subsequent recovery in the survey measures, is consistent with enforcement operating dynamically—not only by directly deterring misconduct, but also by helping rebuild trust and reinforce the norms that support future compliance. Although our design does not separately identify this indirect channel, this interpretation connects our evidence to theoretical and experimental work on the evolution of social norms and their interaction with formal institutions and enforcement (e.g., [Acemoglu and Jackson 2015, 2017](#); [Bisin and Verdier 2024](#); [Benabou and Tirole 2026](#); [Lane et al. 2023](#)). Our model focuses on a government-learning problem in which policymakers infer the persistence of integrity norms from noisy signals of misconduct and adjust formal enforcement accordingly.

Taken together, the survey and administrative evidence support a coherent interpretation of the enforcement episode. The survey evidence indicates that the broader integrity environment deteriorated around the 2011–2012 tightening, while our transaction-level evidence shows that this deterioration coincided with meaningful underlying misconduct that subsequently declined following stronger formal enforcement. While we do not directly measure the persistence of norms, the joint evidence is more consistent with policymakers learning that norms were less persistent than previously believed than with responding to noisy signals alone. More broadly, the findings suggest that even after decades of successful anti-corruption policy, norms alone may not have been sufficiently persistent to sustain a high-integrity equilibrium without continued formal enforcement.

VII. Conclusion

One of the central questions in the study of institutions, social norms, and multiple compliance equilibria is whether decades of successful enforcement can eventually make a high-integrity equilibrium largely self-sustaining, or whether continued formal enforcement remains necessary to preserve compliance. Our evidence suggests that norms alone may not have been sufficiently persistent to sustain a high-integrity equilibrium without continued formal enforcement, even in Singapore, one of the world’s least corrupt countries and a particularly demanding setting in which to study this question.

We reach this conclusion by combining transaction-level evidence on misconduct, survey-based measures of the broader integrity environment, and an enforcement-tightening episode. Civil servants disproportionately purchased homes near planned—but not yet publicly announced—stations, especially among mid-level officials and agencies with greater access to privileged information; related purchases by relatives point to information leakage. Stronger formal enforcement substantially reduced both direct and indirect misconduct. Together with survey evidence showing deterioration and subsequent recovery in perceived corruption, public trust, and pro-integrity norms, these findings

suggest that adverse integrity signals reflected genuine underlying misconduct rather than noise alone.

Our findings have several broader implications beyond our specific context. First, given Singapore's reputation as a high-capacity, low-corruption jurisdiction, our estimates likely provide a useful benchmark, and perhaps a lower bound, on similar effects elsewhere, including in advanced economies. The recent focus on Federal Reserve officials' and legislators' insider trading in the United States illustrates that even mature institutions are not immune to these concerns.

Second, our evidence indicates that norms are fragile: without formal enforcement, ethical behavior can erode over time, even in societies widely regarded as low-corruption and high-norm. This perspective bears directly on current U.S. debates, where perceived weakening of regulatory oversight, including in finance, is viewed by many as cost-saving. Our results caution that reducing enforcement may undermine the very norms that sustain compliance, so the durability of law-abiding behavior cannot be taken for granted.

Third, we believe the framework developed here may be useful beyond the Singapore setting. Governments typically observe noisy signals about the integrity environment—such as perceived corruption, public trust, complaints, or uncovered cases—rather than the underlying misconduct itself. Whenever researchers are able to recover direct measures of otherwise hidden misconduct, combining these two sources of information can improve inference about whether deterioration in the integrity environment reflects genuine behavioral change or merely noisy signals. Although our setting is unique in providing unusually rich administrative data, we expect this general approach to be applicable in a broader range of institutional settings.

Finally, the results speak to policy design and to theories of multiple compliance equilibria and cultural evolution. Even after decades in what appears to be a low-corruption equilibrium, continued reliance on formal enforcement may be necessary to maintain outcomes. A limitation of our analysis is that, while we cast doubt on the strongest version of the self-sustaining-norms hypothesis in this setting, we do not pin down the magnitude or duration of norm persistence. Quantifying persistence and identifying conditions under which norms become self-sustaining remain important tasks for future research.

References

- Acemoglu, D. and Jackson, M. O. (2015). History, expectations, and leadership in the evolution of social norms. *Review of Economic Studies*, 82(2):423–456.
- Acemoglu, D. and Jackson, M. O. (2017). Social norms and the enforcement of laws. *Journal of the European Economic Association*, 15(2):245–295.
- Acemoglu, D. and Verdier, T. (2000). The choice between corruption and market failures. *American Economic Review*, 90(1):194–211.
- Agarwal, S., Koo, K. M., and Sing, T. F. (2015). Impact of electronic road pricing on real estate prices in singapore. *Journal of Urban Economics*, 90:50–59.
- Agarwal, S. and Qian, W. (2014). Consumption and debt response to unanticipated income shocks: Evidence from a natural experiment in singapore. *American Economic Review*, 104(12):4205–4230.
- Agarwal, S., Qian, W., Seru, A., and Zhang, J. (2020). Disguised corruption: Evidence from consumer credit in China. *Journal of Financial Economics*, 137(2):430–450.
- Agarwal, S., Rengarajan, S., Sing, T. F., and Yang, Y. (2016). School allocation rules and housing prices: A quasi-experiment with school relocation events in singapore. *Regional Science and Urban Economics*, 58:42–56.
- Ahern, K. R. (2017). Information networks: Evidence from illegal insider trading tips. *Journal of Financial Economics*, 125(1):26–47.
- Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500.
- Andirin, V. M., Neggers, Y., Shadmehr, M., and Shapiro, J. M. (2025). Surveillance of repression: Theory and implementation. Working Paper 30167, National Bureau of Economic Research. Revised March 2025.
- Avis, E., Ferraz, C., and Finan, F. (2018). Do government audits reduce corruption? estimating the impacts of exposing corrupt politicians. *Journal of Political Economy*, 126(5):1912–1964.
- Axbard, S. (2025). Convicting corrupt officials: Evidence from randomly assigned cases. *The Review of Economic Studies*, 92(4):2271–2302.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76(2):169–217.
- Belmont, W., Sacerdote, B., Sehgal, R., and Van Hoek, I. (2022). Do senators and house members beat the stock market? evidence from the stock act. *Journal of Public Economics*, 207:104602.
- Benabou, R. and Tirole, J. (2026). Laws and norms. *Journal of Political Economy*, 134(2):000–000.
- Bernile, G., Hu, J., and Tang, Y. (2016). Can information be locked up? informed trading ahead of macro-news announcements. *Journal of Financial Economics*, 121(3):496–520.

- Bertrand, M., Bombardini, M., Fisman, R., and Trebbi, F. (2020). Tax-exempt lobbying: Corporate philanthropy as a tool for political influence. *American Economic Review*, 110(7):2065–2102.
- Bertrand, M., Bombardini, M., and Trebbi, F. (2014). Is it whom you know or what you know? an empirical assessment of the lobbying process. *American Economic Review*, 104(12):3885–3920.
- Bidner, C. and Francois, P. (2013). The emergence of political accountability. *Quarterly Journal of Economics*, 128(3):1397–1448.
- Bisin, A. and Verdier, T. (2024). On the joint evolution of culture and political institutions: Elites and civil society. *Journal of Political Economy*, 132(5):1485–1564.
- Bobonis, G. J., Cámara Fuertes, L. R., and Schwabe, R. (2016). Monitoring corruptible politicians. *American Economic Review*, 106(8):2371–2405.
- Cheng, I.-H., Raina, S., and Xiong, W. (2014). Wall street and the housing bubble. *American Economic Review*, 104(9):2797–2829.
- Colonnelli, E. and Prem, M. (2022). Corruption and firms. *The Review of Economic Studies*, 89(2):695–732.
- Djankov, S., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2003). Courts. *The Quarterly Journal of Economics*, 118(2):453–517.
- Egan, M., Matvos, G., and Seru, A. (2017). The market for financial adviser misconduct. *Journal of Political Economy*, 125(4):233–295.
- Eggers, A. and Hainmueller, J. (2013). Capitol losses: The mediocre performance of congressional stock portfolios. *The Journal of Politics*, 75(2):535–551.
- Fisman, R. and Miguel, E. (2007). Corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets. *Journal of Political Economy*, 115(6):1020–1048.
- Giglio, S., Maggiori, M., and Stroebe, J. (2015). Very long-run discount rates. *Quarterly Journal of Economics*, 130(1):1–53.
- Goldsmith-Pinkham, P. and Shue, K. (2023). The gender gap in housing returns. *The Journal of Finance*, 78(2):1097–1145.
- Griffin, J. M. and Maturana, G. (2016). Who facilitated misreporting in securitized loans? *Review of Financial Studies*, 29(2):384–419.
- Guiso, L., Sapienza, P., and Zingales, L. (2004). The role of social capital in financial development. *American Economic Review*, 94(3):526–556.
- Guiso, L., Sapienza, P., and Zingales, L. (2006). Does culture affect economic outcomes? *Journal of Economic Perspectives*, 20(2):23–48.
- Guiso, L., Sapienza, P., and Zingales, L. (2009). Cultural biases in economic exchange? *The Quarterly Journal of Economics*, 124(3):1095–1131.

- He, J., Liu, H., Sing, T. F., Song, C., and Wong, W.-K. (2019). Superstition, conspicuous spending, and housing market: Evidence from singapore. *Management Science*, 66(2):783–804.
- Hu, G. X., Pan, J., and Wang, J. (2017). Early peek advantage? efficient price discovery with tiered information disclosure. *Journal of Financial Economics*, 126(2):399–421.
- Inglehart, R., Haerpfer, C., Moreno, A., Welzel, C., Kizilova, K., Diez-Medrano, J., Lagos, M., Norris, P., Ponarin, E., and Puranen, B. (2022). World values survey: All rounds – country-pooled datafile.
- Irvine, P., Lipson, M., and Puckett, A. (2007). Tipping. *Review of Financial Studies*, 20(3):741–768.
- Kandori, M. (1992). Social norms and community enforcement. *Review of Economic Studies*, 59(1):63–80.
- Khoo, W. (2021). Integrity and independence of public and law enforcement officials in singapore. https://www.unafei.or.jp/publications/pdf/GG14/22_GG14_CP12_Singapore.pdf. Corrupt Practices Investigation Bureau Report.
- La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. (1999). The quality of government. *Journal of Law, Economics, and Organization*, 15(1):222–279.
- La Porta, R., Lopez-de Silanes, F., Shleifer, A., and Vishny, R. W. (1998). Law and finance. *Journal of Political Economy*, 106(6):1113–1155.
- Lane, T., Nosenzo, D., and Sonderegger, S. (2023). Law and norms: Empirical evidence. *American Economic Review*, 113(5):1255–1293.
- Lee, H. L. (2012). Speech by prime minister lee hsien loong at cpib’s 60th anniversary celebration. CPIB. Available at <https://www.cpib.gov.sg/press-room/speeches/speech-prime-minister-lee-hsien-loong-cpibs-60th-anniversary-celebration/>.
- Michaelides, A., Milidonis, A., Nishiotis, G. P., and Papakyriakou, P. (2015). The adverse effects of systematic leakage ahead of official sovereign debt rating announcements. *Journal of Financial Economics*, 116(3):526–547.
- OECD (2013). *Government at a Glance 2013*. OECD Publishing, Paris.
- Olken, B. A. (2007). Monitoring corruption: Evidence from a field experiment in indonesia. *Journal of Political Economy*, 115(2):200–249.
- Olken, B. A. (2009). Corruption perceptions vs. corruption reality. *Journal of Public Economics*, 93(7-8):950–964.
- Piskorski, T., Seru, A., and Witkin, J. (2015). Asset quality misrepresentation by financial intermediaries: Evidence from the rmbs market. *Journal of Finance*, 70(6):2635–2678.
- Posner, E. A. (2002). *Laws and social norms*. Harvard University Press, Cambridge.
- Posner, R. A. (1997). Social norms and the law: An economic approach. *American Economic Review*, 87(2):365–369.

- Quah, J. S. (2020). Combating corruption in singapore: a comparative analysis of two scandals. *Public Administration and Policy*, 23(1):87–99.
- Rajan, R. G. and Zingales, L. (1998). Financial dependence and growth. *The American Economic Review*, 88(3):559.
- Rajan, R. G. and Zingales, L. (2003). The great reversals: the politics of financial development in the twentieth century. *Journal of Financial Economics*, 69(1):5–50.
- Reeb, D. M., Zhang, Y., and Zhao, W. (2014). Insider trading in supervised industries. *Journal of Law and Economics*, 57(3):529–559.
- Rogers, J. L., Skinner, D. J., and Zechman, S. L. C. (2017). Run edgar run: Sec dissemination in a high-frequency world. *Journal of Accounting Research*, 55(2):459–505.
- Shleifer, A. and Vishny, R. W. (1993). Corruption. *The Quarterly Journal of Economics*, 108(3):599–617.
- Shleifer, A. and Vishny, R. W. (1997). A survey of corporate governance. *The Journal of Finance*, 52(2):737–783.
- Shleifer, A. and Vishny, R. W. (1998). *The grabbing hand: Government pathologies and their cures*. Harvard University Press.
- Singapore Department of Statistics (2021). Singapore Census of Population 2020, Statistical Release 2: Households, Geographic Distribution, Transport and Difficulty in Basic Activities. Chapter 7, Transport; Chart 7.1.
- Stephan, A. P., Walther, B. R., and Wellman, L. A. (2021). Profiting from connections: Do politicians receive stock tips from brokerage houses? *Journal of Accounting and Economics*, 72(1):101401.
- Tabellini, G. (2008). The scope of cooperation: Values and incentives. *Quarterly Journal of Economics*, 123(3):905–950.
- Teo, E. (2001). The singapore public service: A development-oriented promotion system. Presentation at the Conference on Career Development of Public Servants in the 21st Century.
- Tirole, J. (1996). A theory of collective reputations (with applications to the persistence of corruption in the firm quality). *Review of Economic Studies*, 63(1):1–22.
- Ziobrowski, A. J., Boyd, J. W., Cheng, P., and Ziobrowski, B. J. (2011). Abnormal returns from the common stock investments of members of the us house of representatives. *Business and Politics*, 13(1):1–22.
- Ziobrowski, A. J., Cheng, P., Boyd, J. W., and Ziobrowski, B. J. (2004). Abnormal returns from the common stock investments of the us senate. *Journal of Financial and Quantitative Analysis*, 39(4):661–676.

**Table 1. Target Property Purchase Intensity Around Announcements:
Descriptive Statistics and Baseline Estimates**

Panels (a)–(b) report descriptive statistics for civil servants and their matched controls, with standard deviations in parentheses. Panel (a) reports observable characteristics in the matched sample, including average age in 2020, approximated monthly income, and the shares of individuals who are male, married, Chinese, Malay, and Indian. Panel (b) reports annual purchase intensity of targeted properties before and after MRT station announcements for the full matched sample using 1 km station radius. Purchase intensity is measured as the number of targeted-property purchases by an individual in an event year, multiplied by 100. The reported difference is Civil Servants minus Matched Control. Panel (c) reports estimated pattern of target-property purchase intensity by civil servants relative to matched non-civil servants around the announcement of new railway station locations. The dependent variable is the number of nearby-property purchases by individual i in event year t , multiplied by 100. Gov is an indicator for civil servant status. $D(t = k)$ denotes event time relative to the announcement year, with $t = -4$ omitted. Columns (1)–(2), (3)–(4), and (5)–(6) define targeted properties using 1 km, 1.5 km, and 2 km station radius, respectively. All specifications pool MRT events and use a balanced event-year panel from four years before to three years after the announcement. Demographic controls include age, gender, marital status, a low-income proxy based on property type of residence, and ethnic-group indicators. Event-year and postal-code fixed effects are included as indicated. Standard errors are clustered by district. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel (a): Observables (Matched Sample)

Variable	Civil Servants	Matched Control	Difference
Age	46.917 (11.620)	47.126 (11.973)	-0.209 (5.186)
Male	0.520 (0.500)	0.507 (0.500)	0.014 (0.445)
Married	0.259 (0.438)	0.244 (0.429)	0.016 (0.352)
Income (K SGD)	5.830 (4.945)	5.418 (4.407)	0.412 (4.350)
Chinese	0.888 (0.316)	0.920 (0.272)	-0.032 (0.319)
Malay	0.055 (0.229)	0.035 (0.184)	0.020 (0.215)
Indian	0.040 (0.196)	0.031 (0.173)	0.009 (0.219)
N	28,529	28,529	–

Panel (b): Annual Purchase Intensity of Targeted Properties

Event Window	Civil Servants	Matched Control	Difference
Pre-announcement	1.404 (0.073)	0.872 (0.057)	0.532 (0.093)
Post-announcement	1.014 (0.057)	0.687 (0.046)	0.326 (0.073)

**Table 1. Target Property Purchase Intensity Around Announcements:
Descriptive Statistics and Baseline Estimates (Continued)**

Panel (c): Baseline Estimates

	Distance $\leq$ 1 km		Distance $\leq$ 1.5 km		Distance $\leq$ 2 km	
	(1)	(2)	(3)	(4)	(5)	(6)
Gov	0.171 (0.194)	0.091 (0.190)	0.374* (0.220)	0.212 (0.203)	0.995*** (0.345)	0.692** (0.276)
Gov×D(t=-3)	0.235 (0.199)	0.235 (0.199)	0.225 (0.221)	0.225 (0.221)	0.096 (0.328)	0.096 (0.328)
Gov×D(t=-2)	0.535** (0.236)	0.535** (0.236)	0.642** (0.268)	0.642** (0.268)	0.364 (0.371)	0.364 (0.371)
Gov×D(t=-1)	0.674** (0.289)	0.674** (0.289)	0.888*** (0.319)	0.888*** (0.319)	0.749** (0.340)	0.749** (0.340)
Gov×D(t=0)	0.193 (0.234)	0.193 (0.234)	0.086 (0.255)	0.086 (0.255)	-0.225 (0.288)	-0.225 (0.288)
Gov×D(t=1)	0.203 (0.218)	0.203 (0.218)	0.492* (0.288)	0.492* (0.288)	0.428 (0.399)	0.428 (0.399)
Gov×D(t=2)	0.000 (0.224)	0.000 (0.224)	0.246 (0.298)	0.246 (0.298)	0.139 (0.420)	0.139 (0.420)
Gov×D(t=3)	0.225 (0.217)	0.225 (0.217)	0.332 (0.231)	0.332 (0.231)	0.021 (0.329)	0.021 (0.329)
Event Year FE	X	X	X	X	X	X
Demographics		X		X		X
Postal Code FE		X		X		X
N	149,552	149,552	149,552	149,552	149,552	149,552
Adj. R-square	0.001	0.099	0.001	0.147	0.001	0.210

Table 2. Heterogeneity — Target Purchase Intensity by Rank and Demographics

The table reports the heterogeneity in the informed purchases across civil servants. The dependent variable is the number of nearby-property (≤ 1 km) purchases by individual i in event year t , multiplied by 100. *Gov* is the dummy variable indicating that the individual is a civil servant. *Gov-High*, *Gov-Mid*, and *Gov-Low* are dummy variables indicating the rank of civil servant. *Gov-Planning* (*Gov-Others*) is the dummy variable for the civil servants working in the government agencies that are (not) responsible for the planning of rail line system. $D(t = k)$ is a time dummy variable that has a value of one if the time is within the k th year before the announcement, or $[k, k+1]$. $D(t=-2$ or $-1)$ is a time dummy that has a value of one if the time is within two years before the announcement. For control variables, we include age, gender, marital status, property type of residence and ethnic groups. Event year and postal code fixed effects are included, and standard errors are clustered by districts. ***, **, and * correspond to statistical significance at 1%, 5% and 10% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Gov	0.095 (0.190)	0.090 (0.190)	0.094 (0.190)	0.093 (0.190)	0.089 (0.190)
Gov $\times D(t=-3)$	0.235 (0.199)	0.235 (0.199)	0.235 (0.199)	0.235 (0.199)	0.235 (0.199)
Gov-High $\times D(t=-2/-1)$	3.422 (3.227)				
Gov-Mid $\times D(t=-2/-1)$	1.343*** (0.314)				
Gov-Low $\times D(t=-2/-1)$	0.390 (0.253)				
Gov-Planning $\times D(t=-2/-1)$		0.807*** (0.270)			
Gov-Others $\times D(t=-2/-1)$		0.378 (0.253)			
Gov $\times D(t=-2/-1)\times$ Old			1.220*** (0.196)		
Gov $\times D(t=-2/-1)\times$ Male				0.520** (0.228)	
Gov $\times D(t=-2/-1)\times$ Chinese					0.392 (0.685)
Gov $\times D(t=-2/-1)\times$ Malay					-0.435 (0.667)
Gov $\times D(t=-2/-1)\times$ Indian					-0.026 (0.877)
Gov $\times D(t=-2/-1)$			-0.394** (0.183)	0.326 (0.233)	0.276 (0.664)
Gov $\times D(t=0)$	0.193 (0.234)	0.193 (0.234)	0.193 (0.234)	0.193 (0.234)	0.193 (0.234)
Gov $\times D(t=1)$	0.203 (0.218)	0.203 (0.218)	0.203 (0.218)	0.203 (0.218)	0.203 (0.218)
Gov $\times D(t=2)$	0.000 (0.224)	-0.000 (0.224)	-0.000 (0.224)	0.000 (0.224)	-0.000 (0.224)
Gov $\times D(t=3)$	0.225 (0.217)	0.225 (0.217)	0.225 (0.217)	0.225 (0.217)	0.225 (0.217)
Event Year FE	X	X	X	X	X
Demographics	X	X	X	X	X
Postal Code FE	X	X	X	X	X
<i>N</i>	149,552	149,552	149,552	149,552	149,552
Adj. R-square	0.099	0.099	0.099	0.099	0.099

Table 3. Cost of Acquiring Target Properties by Civil Servants

The table studies the cost of informed purchase by examining whether civil servants pay a higher price premium for the purchase of target property (defined as those within a radius of 2 km to the station) before the announcement of new railway lines. The dependent variables include total price and price per square meter in logs. *Gov* is an indicator variable that equals one if the transaction is made by civil servants. The sample is a matched sample of civil servants and general population. The housing controls are log(size), even-numbered floor, first floor, freehold, condo. Following He et al. (2019), we also include dummy for floor and unit ending in number 4 and 8 to control for Chinese superstition preference. Transaction Year and spatial fixed effects are also included. Please refer to Appendix A for all variable definitions. Robust standard errors clustered by districts are reported in brackets. ***, **, and * correspond to statistical significance at 1%, 5% and 10% level, respectively.

	Log(Total Price)		Log(Price per square meter)	
	(1)	(2)	(3)	(4)
Gov	-0.077* (0.042)	-0.013 (0.013)	-0.017 (0.032)	-0.013 (0.013)
Gov×D(<i>t</i> = −3)	0.026 (0.039)	0.010 (0.019)	-0.000 (0.038)	0.010 (0.019)
Gov×D(<i>t</i> = −2/−1)	0.085** (0.040)	0.024 (0.018)	0.008 (0.034)	0.024 (0.018)
Gov×D(<i>t</i> = 0)	0.056 (0.044)	0.005 (0.021)	-0.003 (0.039)	0.005 (0.021)
Gov×D(<i>t</i> = 1)	0.104** (0.049)	0.006 (0.023)	0.001 (0.052)	0.006 (0.023)
Gov×D(<i>t</i> = 2)	0.112** (0.052)	0.023 (0.020)	0.012 (0.043)	0.023 (0.020)
Gov×D(<i>t</i> = 3)	0.132** (0.062)	-0.008 (0.026)	0.022 (0.054)	-0.008 (0.026)
Housing Controls		X		X
Demographic Controls		X		X
Spatial FE		X		X
Transaction Year FE	X	X	X	X
N	7,724	7,399	7,438	7,399
Adj. R-square	0.097	0.767	0.363	0.712

Table 4. Comparing Sellers and Buyers

This table compares the sample mean and difference of demographic characteristics for the sellers and buyers in the resale transactions, including age in 2020, approximated monthly income, and the shares of individuals who are male, married, Chinese, Malay, and Indian. Income is approximated using the mean or median income computed by the first four digits of residential postal codes. Panel (a) reports the results for the full sample and panel (b) restricts the sample to transactions that involve affected property at $(t - 1)$ and $(t - 2)$, where t is the corresponding announcement time.

Panel (a): Full Sample

	Buyer		Seller		Difference (Buyer minus Seller)	<i>p</i> -value
	Mean	Sd	Mean	Sd		
Age	54.801	11.542	59.151	12.183	-4.349	0.000
Male (%)	0.524	0.499	0.506	0.500	0.017	0.233
Married (%)	0.360	0.480	0.345	0.476	0.014	0.310
Income Mean (K SGD)	11.383	7.197	11.232	6.375	0.151	0.231
Income Median (K SGD)	11.362	7.339	11.203	6.453	0.159	0.214
Chinese (%)	0.934	0.248	0.940	0.237	-0.006	0.414
Malay (%)	0.022	0.145	0.014	0.118	0.007	0.080
Indian (%)	0.029	0.169	0.028	0.166	0.001	0.852

Panel (b): Affected Property Transactions at $t - 2$ or $t - 1$

	Buyer		Seller		Difference (Buyer minus Seller)	<i>p</i> -value
	Mean	Sd	Mean	Sd		
Age	55.202	11.495	59.185	12.425	-3.982	0.000
Male (%)	0.528	0.499	0.513	0.501	0.016	0.572
Married (%)	0.368	0.482	0.339	0.474	0.029	0.290
Income Mean (K SGD)	11.362	6.779	11.355	5.707	0.008	0.971
Income Median (K SGD)	11.391	7.064	11.345	5.776	0.046	0.834
Chinese (%)	0.926	0.262	0.938	0.241	-0.013	0.384
Malay (%)	0.027	0.161	0.011	0.105	0.015	0.079
Indian (%)	0.034	0.180	0.028	0.165	0.005	0.583


Table 5. Location Manipulation

The table examines whether property purchases by public officers prior to the announcement can ex ante predict the location of new MRT stations. The dependent variable is the distance (kilometers) from the transacted property to the nearest new MRT station. *Gov* equals one when the transaction is by a civil servant. All specifications include transaction-year and district (spatial) fixed effects. Columns (2) and (4) additionally include housing controls: log(size), even-numbered floor, first-floor, freehold, condo, CPF usage, and personal demographics. Robust standard errors clustered by districts are in brackets. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.


	Dependent Variable = Distance to the nearest MRT (in kilometers)			
	(1)	(2)	(3)	(4)
	Purchase prior to the announcement		Purchase post announcement	
Gov	-0.015 (0.021)	-0.010 (0.016)	-0.005 (0.017)	0.006 (0.015)
Demographic Controls		X		X
Housing Controls		X		X
Spatial FE		X		X
Transaction Year FE	X	X	X	X
N	4,289	3,945	3,946	3,706
Adj. R-square	0.018	0.289	0.014	0.270

Figure 1. Map of Mass Rapid Transit (MRT) System

Panel (a) shows the existing and proposed Mass Rapid Transit (MRT) rail lines in Singapore as of March 2020 that are included in the study. Panel (b) plots the yearly evolution of new MRT stations over time. Source: <https://landtransportguru.net/lta-future-mrt-system-map-march-2020>.



Panel (a): Map of Mass Rapid Transit (MRT) System



Panel (b): Evolution of New MRT Stations Over Time


Figure 2. Illustration of MRT Station and Surrounding Residential Area

The figure showcases the geographic distribution of residential properties surrounding Kembangan Station. The red circles indicate, respectively, properties located within a 1 km and 2 km radius of the station.




Figure 3. Informed Trading by Civil Servants Around the Announcement Date

The figure presents the estimated coefficients from the interaction between event-year indicators and the civil servant indicator in the main specification (6), along with two-tailed 95% confidence intervals. The dependent variable is the number of nearby-property purchases by individual i in event year t , multiplied by 100. Point estimates are shown for three different distance cutoffs used to define target properties. The coefficient magnitudes can be interpreted approximately as percentage-point changes in annual nearby-property purchase intensity. Target properties are defined by circular buffers centered on each proposed station, with radius $r = 1, 1.5, 2.0$ km. Event year $t = -4$ is the omitted reference period. The vertical dashed line marks the project announcement date.




**Figure 4. Falsification Tests:
Randomizing Civil Servant Assignment and MRT Locations/Timing**

This figure plots null distributions of $\beta_{-1,-2}$ based on 1,000 bootstrap simulations for each of two falsification exercises. Panel (a) randomizes civil servant status: we randomly assign civil servant status to all individuals in the population dataset during the sample period. Panel (b) jointly randomizes MRT station locations and announcement timing: for each simulation, we redraw MRT station locations within Singapore, assign placebo announcement months, and rebuild the event panel. In each simulation, we estimate the main specification (6), where we recode the event-year dummy by combining two event years, $D(t = -2 \text{ or } -1)$, into a new dummy for $\beta_{-1,-2}$ and re-estimate specification (6) to obtain the $\beta_{-1,-2}$ coefficient estimate. The vertical line indicates the position of the coefficient estimated from the real data, with a highly significant p-value of 0.005.



Panel (a): Randomized Civil Servant Placebo Assignment

**Figure 4. Falsification Tests:
Randomizing Civil Servant Assignment and MRT Locations/Timing
(Continued)**



Panel (b): Randomized MRT Location and Announcement Timing

Figure 5. Placebo Test: Target Purchase Intensity Among HDB Flats

The figure displays coefficients from the interaction between event-year dummies and the civil servant indicator in the main specification (6), along with 95% confidence intervals. The estimates are based on Housing and Development Board (HDB) purchases, with the dependent variable defined as the purchase intensity of target HDB properties. Point estimates are shown for three different distance cutoffs used to define target properties, using the same y-axis scale as Figure 3. For comparison, estimates from Figure 3 are shown in gray in the background. The coefficient magnitudes can be interpreted approximately as percentage-point changes in annual nearby-property purchase intensity. Target properties are defined by circular buffers centered on each proposed station, with radius $r = 1, 1.5, 2.0$ km. Event year $t = -4$ is the omitted reference period. The vertical dashed line indicates the announcement date.






Figure 6. Informed Trading Surrounding the Announcement: Rank and Government Agencies

The figure examines the heterogeneous impact of trading intensity in target properties by civil servants across different ranks and agencies. We re-estimate the main specification (6) after recoding the event-year dummy by combining $t = -2$ and $t = -1$ into a single dummy $D(t = -2 \text{ or } -1)$. The coefficient on this dummy is $\beta_{-1,-2}$. Panel (a) plots the $\beta_{-1,-2}$ coefficient estimates and 95% confidence intervals for the interaction between event-year dummies ($t=-2$ and $t=-1$) and employment rank (low, middle, high), separately for the full sample, MRT station announcements in or before 2011, and MRT station announcements after 2011. Panel (b) presents a similar analysis by government agency type for the same three announcement-year samples, comparing planning departments to other departments. Panel (c) further disaggregates the interactions by employment rank, distinguishing between agency types. The coefficient magnitudes can be interpreted approximately as percentage-point changes in annual nearby-property purchase intensity. All panels use a 1 km distance cutoff to define nearby properties.




Panel (a): Coefficient $\beta_{-1,-2}$ Estimates Across Ranks



Panel (b): Coefficient $\beta_{-1,-2}$ Estimates: Planning vs. Other (Non-Planning) Departments

**Figure 6. Informed Trading Surrounding the Announcement: Rank and Government Agencies
(Continued)**



Panel (c): Coefficient $\beta_{-1,-2}$ Estimates Across Ranks and Agencies

**Figure 7. Purchase Intensity Surrounding the Announcement:
Real Estate Agents and Corporate Board Members**

The figure reports event-study coefficients and 95% confidence intervals from separate regressions for real estate agents and corporate board members, each analogous to (6). Each group is compared with its own matched control sample constructed using the procedure in Section III.C. The dependent variable is the number of purchases by individual i in event year t of target properties located within 1 km of a proposed MRT station, multiplied by 100. Blue circles denote estimates for real estate agents and orange triangles denote estimates for corporate board members. The coefficients can be interpreted approximately as percentage-point changes in the difference in annual nearby-property purchase intensity between each group and its matched controls relative to the omitted reference period, $t = -4$. The vertical dashed line marks the announcement date.







Figure 8. Distribution of Annualized Realized Returns on Informed Purchases

The figure presents summary statistics for the annualized realized returns on informed trades made by civil servants. Levered returns are calculated using the formula in Section V.A, assuming an initial LTV of 80%. Panel (a) shows the raw unlevered (in blue) and levered (in orange) annualized returns at various percentiles of the return distribution for informed purchases by public officers. Panel (b) presents the returns above market returns by calculating the value-weighted market return of all repeated sales that occurred in the same purchase and sale years. Panel (c) further compares these returns by using the value-weighted return of all repeated sales that took place in the same purchase and sale years as the target property purchases by civil servants, but outside the target regions. In Panels (b) and (c), we focus on the difference between the returns on target properties purchased by civil servants and the market returns, as calculated in the manner described above.



Panel (b): Difference in Annualized Realized Unlevered and Levered Returns: Informed Purchases vs. Overall Market


**Figure 8. Distribution of Annualized Realized Returns on Informed Purchases
(Continued)**




Panel (c): Difference in Annualized Realized Unlevered and Levered Returns: Informed Purchases vs. Non-Target Regions

Figure 9. Do Informed Purchases Lead to Faster Sales?

Panel (a) presents summary statistics for the holding period (in months) of two types of housing transactions. The blue bar represents informed purchases made by civil servants that resulted in a subsequent sale. The orange bar shows the distribution of the average holding period for benchmarked transactions, which include all repeated sales occurring at the same purchase time in non-target regions that also led to a subsequent sale. Panel (b) plots the Kaplan-Meier survival curve to track the probability of houses remaining unsold (“surviving”) between informed purchases (by civil servants) and uninformed purchases (occurring at the same time in non-target regions) over the ten-year period since property purchase.




Panel (a): Distribution of Holding Period (in Months)




Panel (b): Kaplan-Meier Survival Curve for House Sale Ex Post

Figure 10. Informed Purchases by Relatives

Panel (a) presents the coefficients on the interaction between the event-year dummies and the relative dummy from the main specification (6), along with 95% confidence intervals. The dependent variable is the purchase intensity of target properties. Event year $t = -4$ is the omitted reference period. The coefficient magnitudes can be interpreted approximately as percentage-point changes in annual nearby-property purchase intensity. The vertical dashed line indicates the announcement date. For panel (a), the control group is matched non-civil-servant buyers without identified ties to civil servants, using the same procedure in Section III.C. Panel (b) presents summary statistics for the annualized realized returns on informed trades made by relatives of civil servants. It shows raw unlevered (in blue) and levered (in orange) annualized returns at various percentiles of the return distribution for informed purchases by relatives of civil servants. Levered returns are calculated using the formula in Section V.A, assuming an initial LTV of 80%.




Panel (a): Target Property Purchases by Relatives/Friends of Civil Servants




Panel (b): Annualized Realized Unlevered and Levered Returns (Unadjusted)

**Figure 11. Role of Regulation:
Informed Purchases Before and After Enforcement Tightening**

Panel (a) presents the coefficients on the interaction between the event-year dummies and the civil servant dummy from the main specification (6), along with 95% confidence intervals. The dependent variable is the number of nearby-property (≤ 1 km) purchases by individual i in event year t , multiplied by 100. Point estimates are shown separately for the earlier (pre-2011 station announcements) and later (post-2011 announcements) sample periods. For panel (a), the control group is a one-to-one matched sample of non-civil-servant buyers, as discussed in Section III.C. Panel (b) shows the corresponding results for purchases by relatives or friends of civil servants. For panel (b), the control group is matched non-civil-servant buyers without identified ties to civil servants, using the same procedure in Section III.C. Event year $t = -4$ is the omitted reference period. The coefficient magnitudes can be interpreted approximately as percentage-point changes in annual nearby-property purchase intensity. The vertical dashed line indicates the announcement date.




Panel (a): Target Property Purchases by Civil Servants: Before and After Enforcement Tightening



Panel (b): Target Property Purchases by Relatives/Friends of Civil Servants: Before and After Enforcement Tightening


Figure 12. Survey-Based Integrity Environment in Singapore

The figure reports three survey-based integrity signals for Singapore. Panels (a) and (b) use the [Asian Barometer Survey](#); Panel (c) uses the [World Values Survey](#). All shares exclude missing, refusal, don't-know, and other non-valid responses. Panel (a) reports the shares of (i) respondents reporting that they or someone they personally know witnessed an act of corruption or bribe-taking by a politician or government official in the past year, and (ii) respondents reporting that most officials or almost everyone in the national government is involved in corruption or bribe-taking. Panel (b) reports harmonized distrust-leaning shares for public institutions, so higher values indicate lower public trust and weaker public-sector reputation. For 2006–2014, the distrust-leaning share combines respondents choosing “not very much trust” and “none at all” on the four-point trust scale. For 2020, which uses a six-point scale, the harmonized distrust-leaning share combines respondents choosing “distrust somewhat,” “distrust a lot,” and “distrust fully.” The 2010–2014 value is the simple average of the available Wave 3 and Wave 4 percentages for each series. For Panels (a) and (b), 2006 and 2020 use Asian Barometer Waves 2 and 5, while the middle-point (2010–14) series averages available Wave 3 and Wave 4 percentages. Panel (c) plots the share choosing “Never justifiable” for accepting a bribe, cheating on taxes, avoiding a fare on public transport, and claiming government benefits to which one is not entitled; higher values indicate stronger pro-integrity norms. The WVS survey years are 2002, 2012, and 2020.




Panel (a): Corruption Perceptions

**Figure 12. Survey-Based Integrity Environment in Singapore
(Continued)**



Panel (b): Public-Sector Distrust and Reputation



Panel (c): Pro-Integrity Norms

Online Appendix

A. Variable Definition

Variable	Definition
Purchase Intensity	The number of target property transactions by buyer b in group i in year t divided by the number of people in group i who are eligible to make the purchase in year t .
Target Property	Dummy that equals one if the house falls within the X km boundary of the new MRT station.
Gov.	Dummy that equals one if the individual works in an administrative agency of the government.
$D(t = k)$	Dummy that equals one if the time period is within the k -th year around the announcement, i.e., $[k, k + 1]$.
Gov_High	Indicator for high-rank public officers.
Gov_Mid	Indicator for mid-rank public officers.
Gov_Low	Indicator for low-rank public officers.
Planning	Indicator for government agencies involved in the planning of rail-line systems.
Non Planning	Indicator for government agencies <i>not</i> involved in the planning of rail-line systems.
Old	Dummy that equals one if the individual's age is above the median.
Male	Dummy that equals one if the individual is male.
Chinese	Dummy that equals one if the individual is Chinese; zero otherwise.
Malay	Dummy that equals one if the individual is Malay; zero otherwise.
Indian	Dummy that equals one if the individual is Indian; zero otherwise.
Size	Total area (square feet) of the property.
Even-numbered floor	Dummy that equals one if the property is located on an even-numbered floor; zero otherwise.
First floor	Dummy that equals one if the property is on the first floor; zero otherwise.
Freehold	Dummy that equals one if the property has 999-year leasehold or freehold; zero otherwise.
Condo	Dummy that equals one if the property is a condominium; zero otherwise.
CPF	Dummy that equals one if the property is purchased using CPF; zero otherwise.
U4 (unit number ending in 4)	Dummy that equals one if the unit number ends in 4; zero otherwise.
U8 (unit number ending in 8)	Dummy that equals one if the unit number ends in 8; zero otherwise.
FL4 (floor number ending in 4)	Dummy that equals one if the floor number ends in 4; zero otherwise.
FL8 (floor number ending in 8)	Dummy that equals one if the floor number ends in 8; zero otherwise.
MCCY	Ministry of Culture, Community and Youth.
MFA	Ministry of Foreign Affairs.
MHA	Ministry of Home Affairs.
MINDEF	Ministry of Defence.
MND	Ministry of National Development.
MOE	Ministry of Education.
MOF	Ministry of Finance.
MOH	Ministry of Health.
MOM	Ministry of Manpower.
MOT	Ministry of Transport.
MSF	Ministry of Social and Family Development.
MTI	Ministry of Trade and Industry.
PMO	Prime Minister's Office.
MEWR	Ministry of Environment and Water Resources.
MCI	Ministry of Communications and Information.
MINLAW	Ministry of Law.
OOS	Organs of State.

B. Classifying Civil Servant Rank Tier using AI

We classify civil-service positions using an LLM-assisted protocol based on each official's agency and designation. The protocol maps titles into three rank tiers—high, middle, and low—using pre-specified definitions. We then validate the classification on a hand-coded sample. The primary objective is to categorize government personnel in Singapore into distinct rank tiers using the agency and post information in the civil servant data. Here we describe the detailed instruction given to ChatGPT 4.0 for this task. Specifically, we ask ChatGPT to take the role of a data analyst, specializing in analyzing and tagging the most suitable rank group for Singapore civil servants.

Core Prompt.

You are a GPT with a formal and academically rigorous approach, specializing in analyzing and tagging data for Singapore government employees. You process data in Excel or CSV formats, adding detailed labels and descriptions to each entry. Your analysis considers multiple factors for accurate categorization, aiming to provide in-depth, unrestricted results suitable for academic purposes. You communicate in a formal, precise manner, reflecting your academic integrity and the seriousness of your task. Your role is to assist users in obtaining a nuanced, scholarly understanding of government employee data. Here offers the definition and example for High/Middle/Low group.

- (1) **High:** *This category includes the highest-ranking officials in the Civil Service. They are responsible for policy formulation, implementation, and overall resource management.*
 - *Example: Administrative Officer (AO) in the Administrative Service, who are at the top of the hierarchy.*
- (2) **Upper Middle:** *Officers in this category typically hold senior management positions, overseeing significant aspects of public service. This group consists of mid-level management officers who are responsible for managing specific areas within a ministry or department.*
 - *Example: Senior Management Executive Officer (SMEO) in the Administrative Service, handling major administrative tasks and decision-making processes. Middle management levels in the Management Executive Scheme (MX), such as MX10, who oversee specific areas like human resources or finance.*
- (3) **Upper Low:** *Officers in this category provide specialized support or manage smaller teams or units within a larger division. This category also includes clerical and administrative staff who play a crucial role in the day-to-day functioning of the Civil Service.*
 - *Example: Entry-level positions in the Management Executive Scheme (MX), like MX9, who are involved in specific operational tasks within their departments. Clerical and administrative staff in the Support Staff group, providing essential support services for the smooth functioning of their respective offices.*

C. Civil Service Pay Grades

The public service system is generally divided into high, middle, and low levels. At the top level, about 300 public officers make up the elite Administrative Service and occupy key leadership roles in government ministries, major statutory boards, and government-linked companies. The structure of the bureaucratic organization aligns with the Westminster parliamentary form as a career service that is subordinate and loyal to the government (Painter 2004). Ministerial salaries are benchmarked against potential highest earnings in private sector professions (Public Service Division, White Paper on Salaries for a Capable and Committed Government, 2012).³³ For instance, the income benchmark for an M4 Minister (most junior) is 40% less than the median income of the top 1,000 Singapore Citizens earners.³⁴ As of 2024, the annual salary for an entry-level MR4 grade minister was raised to SG\$1.2 million from SG\$1.1 million.

The mid and low levels comprise most of the public service pool. Mid-level officers typically hold senior management positions and oversee significant aspects of the public service. This includes mid-level management officers responsible for managing specific areas within a ministry or department. Low-level officers provide specialized support or manage smaller teams within a larger division, including entry, clerical, and administrative staff. There is no official release of the pay grades for this category, and the reference salary shown below is sourced from unofficial forums.³⁵

Pay Grade	Pay Range
MX9 (Directors)	\$16,700 - \$22,700
MX10 (Senior Management)	\$7,600 - \$12,300
MX11A (Management)	\$5,900 - \$9,700
MX11 (Management)	\$4,800 - \$7,700
MX12 (Executives)	\$3,300 - \$6,250

³³Public Service Division 2012, White Paper on Salaries for a Capable and Committed Government, accessed 8 April 2024, <https://www.psd.gov.sg/files/white-paper---salaries-for-a-capable-and-committed-government.pdf>

³⁴See <https://dollarsandsense.sg/heres-much-singapores-president-cabinet-ministers-paid-salary/>.

³⁵See <https://blog.seedly.sg/civil-service-salary/>.

D. Additional Analysis

D1. Informed Trading Across Rank and Government Agencies: Difference Test

This table assesses whether there are significant heterogeneous impacts in informed trading intensity across ranks and ministries, as presented in Figure 6. Panels (a) and (b) illustrate the difference in estimated coefficient values $\beta_{-1,-2}$ around the announcement between mid- and low-rank officers, and mid- and high-rank officers respectively. Panel (c) illustrates the difference in estimated coefficient values between officers at planning and non-planning agencies. Each cell presents the difference in coefficients and their associated p-values. ***, **, and * correspond to statistical significance at 1%, 5% and 10% level respectively.

Panel (a): Rank (Middle Minus Low)			
	Full Sample	Earlier Events (≤ 2011)	Later Events (> 2011)
Difference	0.953***	0.942***	0.011
(p-value)	(0.003)	(0.003)	(0.891)
Panel (b): Rank (Middle Minus High)			
	Full Sample	Earlier Events (≤ 2011)	Later Events (> 2011)
Difference	-2.079	-2.669	0.590
(p-value)	(0.528)	(0.455)	(0.153)
Panel (c): Agency (Planning Minus Non-Planning)			
	Full Sample	Earlier Events (≤ 2011)	Later Events (> 2011)
Difference	0.429*	0.412*	0.018
(p-value)	(0.069)	(0.057)	(0.760)

D2. Descriptive Statistics for Real Estate Agents and Corporate Board Members

Panels (a)–(b) report descriptive statistics for real estate agents and corporate board members and their respective matched controls, with standard deviations in parentheses. Panel (a) reports observables in the matched samples, including average age and the shares of individuals who are male, married, HDB residents, Chinese, Malay, Indian, or of other ethnicity. HDB resident is an indicator for residence in public housing. Panel (b) reports annual purchase intensity of targeted properties before and after MRT station announcements for the full matched samples using 1 km station radius. Purchase intensity is measured as the number of targeted-property purchases by an individual in an event year, multiplied by 100. The reported difference is Real Estate Agents or Board Members minus Matched Control.

Panel (a): Observables (Matched Samples)

Variable	Real Estate Agents	Matched Control	Difference
Age	47.952 (10.804)	48.071 (11.540)	-0.119 (0.101)
Male	0.564 (0.496)	0.542 (0.498)	0.022 (0.004)
Married	0.247 (0.431)	0.221 (0.415)	0.026 (0.004)
HDB resident	0.746 (0.436)	0.786 (0.410)	-0.041 (0.004)
Chinese	0.916 (0.277)	0.939 (0.240)	-0.023 (0.002)
Malay	0.047 (0.211)	0.032 (0.177)	0.015 (0.002)
Indian	0.025 (0.156)	0.020 (0.142)	0.005 (0.001)
Other ethnicity	0.012 (0.109)	0.008 (0.092)	0.004 (0.001)
N	24,493	24,493	–

Variable	Board Members	Matched Control	Difference
Age	66.796 (9.323)	66.515 (10.066)	0.282 (0.371)
Male	0.918 (0.275)	0.915 (0.279)	0.003 (0.011)
Married	0.597 (0.491)	0.614 (0.487)	-0.018 (0.019)
HDB resident	0.136 (0.343)	0.337 (0.473)	-0.200 (0.016)
Chinese	0.949 (0.219)	0.955 (0.208)	-0.005 (0.008)
Malay	0.014 (0.117)	0.011 (0.104)	0.003 (0.004)
Indian	0.023 (0.151)	0.026 (0.158)	-0.002 (0.006)
Other ethnicity	0.013 (0.114)	0.009 (0.093)	0.004 (0.004)
N	1,364	1,364	–

Panel (b): Annual Purchase Intensity of Targeted Properties

Event Window	Real Estate Agents	Matched Control	Difference
Pre-announcement	1.130 (0.035)	0.866 (0.033)	0.264 (0.048)
Post-announcement	0.991 (0.032)	0.692 (0.029)	0.299 (0.043)

Event Window	Board Members	Matched Control	Difference
Pre-announcement	3.427 (0.265)	3.226 (0.359)	0.202 (0.446)
Post-announcement	2.511 (0.223)	2.071 (0.269)	0.440 (0.349)


D3. Informed Trading by Corporate Board Members: Heterogeneity

The table reports the heterogeneity in the informed purchase across corporate board members. The dependent variable is the number of target property purchases multiplied by 100. *Board-Multi* (*Board-Single*) is the dummy variable indicating whether the board member sits on multiple (or single) boards. *Board-Large* (*Board-Small*) equals one if the market value of the firm that the board member belongs to is above (below) the median value of the distribution. We also include dummies for board members as well as event-year in Table 2. For control variables, we include age, gender, marital status, low-income dummy (proxied by property type of residence) and ethnic groups. Event year and postal code fixed effects are included, and standard errors are clustered by districts. ***, **, and * correspond to statistical significance at 1%, 5% and 10% level respectively.


	(1)	(2)
Board-Multi $\times D(t=-2/-1)$	0.647 (1.000)	
Board-Single $\times D(t=-2/-1)$	0.520 (0.840)	
Board-Large $\times D(t=-2/-1)$		0.907 (0.966)
Board-Small $\times D(t=-2/-1)$		0.113 (0.774)
Event Year FE	X	X
Demographics	X	X
Postal Code FE	X	X
N	21,824	21,824
Adj. R-square	0.109	0.109

D4. Informed Trading by Relatives: Annualized Realized Returns

The figure presents summary statistics for the annualized realized returns on informed trades made by relatives of civil servants. Levered returns are calculated using the formula in Section V.A, assuming an initial LTV of 80%. Panel (a) presents the returns above market returns by calculating the value-weighted market return of all repeated sales that occurred in the same purchase and sale years. Panel (b) further compares these returns by using the value-weighted return of all repeated sales that took place in the same purchase and sale years as the target property purchases by relatives, but outside the target regions. In Panels (a) and (b), we focus on the difference between the returns on target properties purchased by relatives and the market returns, as calculated in the manner described above.




Panel (a): Difference in Annualized Realized Unlevered and Levered Returns: Informed Purchases vs. Overall Market




Panel (b): Difference in Annualized Realized Unlevered and Levered Returns: Informed Purchases vs. Non-Target Regions

D5. Informed Trading by Relatives: Holding Periods

Panel (a) presents summary statistics for the holding period (in months) of two types of housing transactions. The blue bar represents informed purchases made by relatives of civil servants that resulted in a subsequent sale. The orange bar shows the distribution of the average holding period for benchmarked transactions, which include all repeated sales occurring at the same purchase time in non-target regions that also led to a subsequent sale. Panel (b) plots the Kaplan-Meier survival curve to track the probability of houses remaining unsold (“surviving”) between informed purchases (by relatives) and uninformed purchases (occurring at the same time in non-target regions) over the ten-year period since property purchase.



Panel (a): Distribution of Holding Period (in Months)



Panel (b): Hazard Model of House Sale Ex Post

E. Model Solution and Proofs

E1. Equilibrium and Comparative Statics

Individual cutoff and honesty shares. Given uniform enforcement p_t and public belief $\hat{R}_{t-1}$, an official of rank i prefers honesty iff

$$\gamma_i \mathbb{E}_t[\hat{R}_{t-1}] \geq B_i - \theta - p_t F_i + (1 - p_t) \gamma_i \mathbb{E}_t[\hat{R}_{t-1}] - \epsilon_t \iff \theta \geq \theta_{i,t}^* \equiv B_i - p_t (F_i + \gamma_i \mathbb{E}_t[\hat{R}_{t-1}]) - \epsilon_t.$$

With $\theta \sim U[0, 1]$, the honesty share (clipped to $[0, 1]$) is

$$x_{i,t} = [1 - \theta_{i,t}^*]_{[0,1]} = [1 - B_i + p_t (F_i + \gamma_i \mathbb{E}_t[\hat{R}_{t-1}]) + \epsilon_t]_{[0,1]}.$$

On the interior, $\partial x_{i,t} / \partial p_t = F_i + \gamma_i \hat{R}_{t-1}$, $\partial x_{i,t} / \partial \hat{R}_{t-1} = p_t \gamma_i$, and $\partial^2 x_{i,t} / (\partial p_t \partial \hat{R}_{t-1}) = \gamma_i > 0$, confirming the Complementarity between Social Norms and Legal Enforcement.

Optimal enforcement. Given its inability to observe rank-specific parameters, the government is assumed to act as if it faces a single, representative official, effectively aggregating all ranks. The government holds a prior belief that $\psi_t \sim \mathcal{N}(\mu_{\psi_t}, \sigma_{\psi_t}^2)$, where $\psi_t = \ln\left(\frac{\hat{\phi}_t}{1 - \hat{\phi}_t}\right)$, and $\hat{\phi}_t \in (0, 1)$. Given belief $\hat{\phi}_t$, the first-order condition for $\max_{p_t \in [0,1]} \mathbb{E}_t[R_t] - C(p_t)$ within the interior is

$$\mathbb{E}_t \left[(1 - \hat{\phi}_t) \rho \cdot \frac{\partial \hat{x}_t}{\partial p_t} - C'(p_t) \right] = 0.$$

With $C(p) = \frac{\kappa}{2} p^2$ and interior solution,

$$p_t^{in*} = \mathbb{E}_t \left[\frac{(1 - \hat{\phi}_t) \rho (F + \gamma \hat{R}_{t-1})}{\kappa} \right]_{[0,1]}.$$

And the corner solution, in the sense that $\mathbb{E}(x_t) = 1$, is

$$\bar{p}_t = \mathbb{E}_t \left[\frac{B}{F + \gamma \hat{R}_{t-1}} \right]$$

which is decreasing in $\hat{R}_{t-1}$, $\frac{\partial \bar{p}_t}{\partial \hat{R}_{t-1}} < 0$. If κ is low enough, a government can always optimally set enforcement at a level, $\bar{p}_t$, that ensures $x_t = 1$ in expectation. Consequently, over time, R_t is strictly increasing while the optimal enforcement $p_t^* = \bar{p}_t$ is strictly decreasing, if we ignore ν_t . This continues until period τ , when the system reaches its steady state, where $R_\tau = \rho$ and $R_{\tau+1} = \phi R_\tau + (1 - \phi) \rho = R_\tau$. This equilibrium path demonstrates that with a sufficiently low enforcement cost or a strong government preference for social integrity, robust early enforcement can elevate social norms and, in turn, permit lower enforcement in later periods. This confirms that High social norms lower the enforcement floor.

Put together, given that $\hat{R}_{t-1} = \hat{R}_{t-1}(\hat{\phi}_t)$ except for $\hat{R}_0 = \bar{R}_0 = R_0 + \eta_0$, the optimal enforcement is

$$p_t^* = \min \left\{ p_t^{in*} = \mathbb{E}_t \left[\frac{(1 - \hat{\phi}_t) \rho (F + \gamma \hat{R}_{t-1}(\hat{\phi}_t))}{\kappa} \right]_{[0,1]}, \bar{p}_t = \mathbb{E}_t \left[\frac{B}{F + \gamma \hat{R}_{t-1}(\hat{\phi}_t)} \right] \right\}.$$

Belief updating for norm persistence. Using the noisy signal $\tilde{R}_t = R_t + \eta_t$ and the state law $R_t = \phi R_{t-1} + (1 - \phi)\rho x_t + \nu_t$ with $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$, $\nu_t \sim \mathcal{N}(0, \sigma_\nu^2)$, $\eta_t \perp \nu_t$, define the prediction error $e_t \equiv \tilde{R}_t - \hat{R}_t$, where $\hat{R}_t$ is the model-implied prediction given $\hat{x}_t = \hat{x}_t(p_t^*)$. Since the predicted signal $\hat{R}_t$ is a non-linear function of the underlying belief state ψ_t , we use the Extended Kalman Filter (EKF) to update the government's belief. The EKF handles this non-linearity by performing a local linearization of the system at the current point estimate. First, we define the Jacobian of the prediction function with respect to the state ψ_t , evaluated at the prior mean μ_{ψ_t} , $H_t \equiv \left. \frac{\partial \hat{R}_t}{\partial \psi_t} \right|_{\psi_t = \mu_{\psi_t}}$. Next, we compute the optimal Kalman Gain, K_t , which minimizes the posterior error covariance: $K_t \equiv \frac{\sigma_{\psi_t}^2 H_t}{H_t^2 \sigma_{\psi_t}^2 + (\sigma_\eta^2 + \sigma_\nu^2)}$. Finally, the government's prior belief is updated to a posterior belief, $\psi_{t+1} \sim \mathcal{N}(\mu_{\psi_{t+1}}, \sigma_{\psi_{t+1}}^2)$, using the prediction error and the Kalman Gain, we have $\mu_{\psi_{t+1}} = \mu_{\psi_t} + K_t e_t$, $\sigma_{\psi_{t+1}}^2 = (1 - K_t H_t) \sigma_{\psi_t}^2$. This posterior belief then serves as the prior, with $\hat{\phi}_{t+1} = \frac{e^{\mu_{\psi_{t+1}}}}{1 + e^{\mu_{\psi_{t+1}}}}$, for the decision-making process in period $t+1$. In particular, at time t , we can sign H_t under a plausible economic assumption. For a society transitioning from a high-norm period $t-1$, we assume the stock of reputational capital exceeds the anticipated honesty flow, i.e., $\hat{R}_{t-1} > \rho \hat{x}_t$. This condition holds when R_{t-1} is sufficiently high and the private benefits of misconduct, B , are not excessive. The Jacobian H_t , evaluated at the prior mean, is strictly positive:

$$H_t = \left[\hat{R}_{t-1} - \rho \hat{x}_t \right] \hat{\phi}_t (1 - \hat{\phi}_t) > 0,$$

which means that $K_t > 0$ and $K_t H_t > 0$. This updating process has two key implications: (i) *a negative surprise* ($e_t < 0$) *leads the government to revise its belief about norm persistence downward* ($\mu_{\psi_{t+1}} < \mu_{\psi_t}$), and (ii) the posterior belief is more precise than the prior, as its variance is reduced ($\sigma_{\psi_{t+1}}^2 < \sigma_{\psi_t}^2$).

Relationship between optimal enforcement and belief of persistence. When R_t is anchored with an initial norm $\tilde{R}_0$, $\hat{R}_{t-1}(\hat{\phi}_t) = \hat{\phi}_t \hat{R}_{t-2} + (1 - \hat{\phi}_t) \rho \hat{x}_{t-1} + \nu_{t-1} = \hat{\phi}_t^{t-1} \tilde{R}_0 + \rho \sum_{i=1}^{t-1} \hat{\phi}_t^{t-1-i} (1 - \hat{\phi}_t) \hat{x}_i + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} \nu_j$, and $x_t = [1 - B + p_t(F + \gamma \hat{R}_{t-1})]_{[0,1]}$. Below we discuss the relationship between p_t^* and $\hat{\phi}_t$:

- (1) *Minimum enforcement to sustain full honesty and its dependence on persistence.* We first consider a special case, in which the government targets full honesty in period t (i.e., chooses $\bar{p}_t$ such that $x_t = 1$, ignoring ϵ_t and ν_t), the minimum enforcement level is

$$\bar{p}_t = \mathbb{E}_t \left[\frac{B}{F + \gamma \hat{R}_{t-1}} \right], \quad \text{which is strictly decreasing in } \hat{R}_{t-1}.$$

Along a path with strong enforcement in all prior periods so that $x_s = 1$ for $s = 1, \dots, t-1$, reputation evolves as

$$\begin{aligned}\hat{R}_{t-1}(\hat{\phi}_t) &= \hat{\phi}_t^{t-1} \tilde{R}_0 + \rho \sum_{i=1}^{t-1} \hat{\phi}_t^{t-1-i} (1 - \hat{\phi}_t) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j \\ &= \rho + \hat{\phi}_t^{t-1} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j,\end{aligned}$$

hence

$$\bar{p}_t(\hat{\phi}_t) = \mathbb{E}_t \left[\frac{B}{F + \gamma [\rho + \hat{\phi}_t^{t-1} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j]} \right].$$

For $t \geq 2$, the comparative static with respect to the persistence belief $\hat{\phi}_t$ is

$$\frac{\partial \bar{p}_t}{\partial \hat{\phi}_t} = -\mathbb{E}_t \left[\frac{B\gamma \left[(t-1) \hat{\phi}_t^{t-2} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} (t-1-j) \hat{\phi}_t^{t-2-j} v_j \right]}{\left\{ F + \gamma [\rho + \hat{\phi}_t^{t-1} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j] \right\}^2} \right].$$

Therefore, the sign and magnitude of $\partial \bar{p}_t / \partial \hat{\phi}_t$ are governed by how a "memory-weighted" average of past innovations has moved reputation going into period t . Three implications follow: (1) Dominance of recent history (large t). As t grows, the initial-gap term $(t-1) \hat{\phi}_t^{t-2} (\tilde{R}_0 - \rho)$ shrinks rapidly, so the comparative static is driven by the weighted sum of recent-to-intermediate shocks $\{v_j\}$, with weights $(t-1-j) \hat{\phi}_t^{t-2-j}$. Note that the most recent shock v_{t-1} has weight zero. (2) High recent/pre-period norms $\Rightarrow \partial \bar{p}_t / \partial \hat{\phi}_t < 0$. If the memory-weighted sum of shocks is positive (recent good news buoyed reputation), then raising perceived persistence $\hat{\phi}_t$ increases expected $\hat{R}_{t-1}$, so the minimum enforcement needed to keep $x_t = 1$ falls. Intuitively, favorable reputational momentum lets the government "lean on norms." (3) Recent slippage $\Rightarrow \partial \bar{p}_t / \partial \hat{\phi}_t > 0$. If the memory-weighted shocks are negative (recent deterioration), higher $\hat{\phi}_t$ makes bad realizations stickier, lowering expected $\hat{R}_{t-1}$ and raising the minimum enforcement required for full honesty.

Remark (Singapore, late 2000s: high pre-period norm). For the pre-2011 episode, the empirical narrative is that reputation going into t was high relative to ρx (near-full honesty in recent periods). This is consistent with large t , an initial anchor below steady state ($\tilde{R}_0 < \rho$), and a run of positive recent shocks that buoyed reputation. Define the memory-weighted shock index $W_t := (t-1) \hat{\phi}_t^{t-2} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} (t-1-j) \hat{\phi}_t^{t-2-j} v_j$. The derivative above implies $\partial \bar{p}_t / \partial \hat{\phi}_t = -\mathbb{E}_t \left[B\gamma W_t / \{F + \gamma [\rho + \hat{\phi}_t^{t-1} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j]\}^2 \right]$, so $\partial \bar{p}_t / \partial \hat{\phi}_t < 0$ whenever $W_t > 0$. In the late-2000s benchmark, $W_t > 0$ holds under three mild conditions: (i) the initial-gap term $(t-1) \hat{\phi}_t^{t-2} (\tilde{R}_0 - \rho)$ is negative but small for large t and $\hat{\phi}_t \in (0, 1)$ (geometric decay dominates the linear factor); (ii) recent shocks are on net positive so that $\sum_{j=1}^{t-1} (t-1-j) \hat{\phi}_t^{t-2-j} v_j > 0$; and (iii) the interior feasibility condition $F + \gamma \hat{R}_{t-1} > 0$ (and any policy cap satisfies $\bar{p}_t \leq p_{\max}$)

holds. Note that the most recent innovation v_{t-1} has zero weight in W_t ; the effective memory window is governed by geometric decay with approximate length of order $1/(1 - \hat{\phi}_t)$. A useful boundary case is $t = 2$, where $\partial \bar{p}_2 / \partial \hat{\phi}_2 = -\text{sign}(\tilde{R}_0 - \rho)$; for larger t the initial-condition effect fades and the sign is determined by the recent reputation path. Hence, in the Singapore late-2000s setting with high pre-period norms, greater perceived persistence $\hat{\phi}_t$ raises expected $\hat{R}_{t-1}$ and *reduces* the minimum enforcement $\bar{p}_t$ required to sustain $x_t = 1$: the government can lean on norms and economize on enforcement.

(2) *Interior optimal enforcement and its dependence on persistence.* Recall the interior policy

$$p_t^{\text{in}*} = \left[\mathbb{E}_t \left(\frac{(1 - \hat{\phi}_t) \rho (F + \gamma \hat{R}_{t-1})}{\kappa} \right) \right]_{[0,1]}.$$

Throughout this subsection we study the *unclipped* interior solution and then comment on the $[0, 1]$ projection. Define $A_t \equiv F + \gamma \hat{R}_{t-1}$. For interior points,

$$p_t^{\text{in}*} = \frac{\rho}{\kappa} \mathbb{E}_t[(1 - \hat{\phi}_t) A_t].$$

Two elementary partial derivatives (holding the other argument fixed) are:

$$\left. \frac{\partial p_t^{\text{in}*}}{\partial \hat{R}_{t-1}} \right|_{\hat{\phi}_t} = \frac{\rho}{\kappa} (1 - \mathbb{E}_t[\hat{\phi}_t]) \gamma > 0, \quad \left. \frac{\partial p_t^{\text{in}*}}{\partial \hat{\phi}_t} \right|_{\hat{R}_{t-1}} = -\frac{\rho}{\kappa} A_t < 0.$$

Hence enforcement rises with reputational salience and falls with perceived persistence *ceteris paribus*. In general $\hat{R}_{t-1} = \hat{R}_{t-1}(\hat{\phi}_t)$, so

$$\frac{\partial p_t^{\text{in}*}}{\partial \hat{\phi}_t} = \frac{\rho}{\kappa} \mathbb{E}_t \left[-A_t + (1 - \hat{\phi}_t) \gamma \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} \right].$$

The first (negative) term is the *mechanical substitution effect*: more persistence reduces the need for force today. The second (positive) term is the *reputational amplification*: if higher persistence raises the perceived stock $\hat{R}_{t-1}$, it boosts the marginal effectiveness of enforcement. A sufficient condition for $p_t^{\text{in}*}$ to be strictly decreasing in $\hat{\phi}_t$ is

$$(1 - \hat{\phi}_t) \gamma \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} < A_t = F + \gamma \hat{R}_{t-1},$$

i.e., the amplification is dominated by the level of (legal + reputational) marginal deterrence.³⁶

Along histories with strong prior enforcement $x_s = 1$ for $s = 1, \dots, t-1$,

$$\hat{R}_{t-1}(\hat{\phi}_t) = \rho + \hat{\phi}_t^{t-1} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j, \quad \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} = (t-1) \hat{\phi}_t^{t-2} (\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} (t-1-j) \hat{\phi}_t^{t-2-j} v_j.$$

³⁶With heterogeneous ranks, drop ranks at corners from A_t in the local slope.

Substituting yields the explicit slope

$$\begin{aligned} \frac{\partial p_t^{\text{in*}}}{\partial \hat{\phi}_t} = & -\frac{\rho}{\kappa} \left[F + \gamma(\rho + \hat{\phi}_t^{t-1}(\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j) \right] \\ & + \frac{\rho}{\kappa} (1 - \hat{\phi}_t) \gamma \left[(t-1) \hat{\phi}_t^{t-2}(\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} (t-1-j) \hat{\phi}_t^{t-2-j} v_j \right]. \end{aligned}$$

Remark. From $\partial p_t^{\text{in*}}/\partial \hat{\phi}_t = (\rho/\kappa) \mathbb{E}_t[-A_t + (1 - \hat{\phi}_t)\gamma \partial \hat{R}_{t-1}/\partial \hat{\phi}_t]$ with $A_t = F + \gamma \hat{R}_{t-1}$, the first term is a negative mechanical substitution effect (greater persistence substitutes for current enforcement), while the second term is a reputational amplification term whose sign follows $\partial \hat{R}_{t-1}/\partial \hat{\phi}_t = (t-1)\hat{\phi}_t^{t-2}(\tilde{R}_0 - \rho) + \sum_{j=1}^{t-1} (t-1-j)\hat{\phi}_t^{t-2-j} v_j$. In the pre-2011 Singapore setting, norms entering period t are high, so A_t is large; t is large and $\tilde{R}_0 < \rho$ so the initial-gap component is small and negative; and recent memory-weighted shocks are on net positive, implying $\partial \hat{R}_{t-1}/\partial \hat{\phi}_t > 0$. Two forces then tilt the sign towards a decrease: (i) the large A_t makes the negative substitution term quantitatively strong under high norms; and (ii) the amplification term is attenuated by the factor $(1 - \hat{\phi}_t)$, which is small when persistence is perceived to be high. A sufficient condition we need is $(1 - \hat{\phi}_t)\gamma \partial \hat{R}_{t-1}/\partial \hat{\phi}_t < A_t = F + \gamma \hat{R}_{t-1}$, under which $\partial p_t^{\text{in*}}/\partial \hat{\phi}_t < 0$. Intuitively, with elevated reputational capital, marginal deterrence is already strong, so higher persistence mainly reduces the need for force today. If recent reputation had been slipping (memory-weighted shocks negative), then $\partial \hat{R}_{t-1}/\partial \hat{\phi}_t < 0$ and the slope becomes even more negative. Projecting back to $[0, 1]$ does not alter these signs on the interior; on clipped regions the slope is zero until the policy re-enters the interior.

Interior policy with arbitrary history (no $x \equiv 1$ assumption). We keep the interior policy

$$p_t^{\text{in*}} = \frac{\rho}{\kappa} \mathbb{E}_t[(1 - \hat{\phi}_t) A_t], \quad A_t \equiv F + \gamma \hat{R}_{t-1}.$$

Allowing for an arbitrary realized history $\{\hat{x}_s\}_{s \leq t-1}$ and endogenous policy,

$$\hat{R}_{t-1}(\hat{\phi}_t) = \hat{\phi}_t^{t-1} R_0 + \rho \sum_{i=1}^{t-1} (1 - \hat{\phi}_t) \hat{\phi}_t^{t-1-i} \hat{x}_i(\hat{\phi}_t) + \sum_{j=1}^{t-1} \hat{\phi}_t^{t-1-j} v_j.$$

Differentiating yields a decomposition into a *mechanical* term (holding $\hat{x}_i$ fixed) and a *behavioral* term (via endogenous $\hat{x}_i$):

$$\frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} = \underbrace{\sum_{s=1}^{t-1} \hat{\phi}_t^{t-1-s} (\hat{R}_{s-1} - \rho \hat{x}_s)}_{\mathcal{M}_{t-1}(\hat{\phi}_t)} + \underbrace{\rho \sum_{s=1}^{t-1} (1 - \hat{\phi}_t) \hat{\phi}_t^{t-1-s} \frac{\partial \hat{x}_s}{\partial \hat{\phi}_t}}_{\mathcal{B}_{t-1}(\hat{\phi}_t)}.$$

On the interior (no truncation in $\hat{x}_s$), $\hat{x}_s = 1 - B + p_s^{\text{in}*} (F + \gamma \hat{R}_{s-1})$, so

$$\frac{\partial \hat{x}_s}{\partial \hat{\phi}_t} = (F + \gamma \hat{R}_{s-1}) \frac{\partial p_s^{\text{in}*}}{\partial \hat{\phi}_t} + \gamma p_s^{\text{in}*} \frac{\partial \hat{R}_{s-1}}{\partial \hat{\phi}_t}.$$

Substituting into

$$\begin{aligned} \frac{\partial p_t^{\text{in}*}}{\partial \hat{\phi}_t} &= \frac{\rho}{\kappa} \mathbb{E}_t \left[-A_t + (1 - \hat{\phi}_t) \gamma \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} \right] \\ &= \frac{\rho}{\kappa} \mathbb{E}_t \left[-A_t + (1 - \hat{\phi}_t) \gamma \{ \mathcal{M}_{t-1}(\hat{\phi}_t) + \mathcal{B}_{t-1}(\hat{\phi}_t) \} \right] \end{aligned}$$

gives the general comparative static without assuming $x \equiv 1$, where a negative substitution force $-A_t$ is opposed by an amplification force $(1 - \hat{\phi}_t) \gamma \{ \mathcal{M}_{t-1} + \mathcal{B}_{t-1} \}$ that operates through how higher persistence raises the inferred reputational stock. The mechanical component $\mathcal{M}_{t-1}$ aggregates history through weights $\hat{\phi}_t^{t-1-s}$ and is positive when past periods looked “tight and clean” in the sense $\hat{R}_{s-1} \gtrsim \rho \hat{x}_s$. The behavioral component $\mathcal{B}_{t-1}$ captures how reinterpreting past honesty given a higher $\hat{\phi}_t$ shifts inferred $\hat{x}_s$; on the interior it is nonnegative whenever $\partial \hat{R}_{s-1} / \partial \hat{\phi}_t \geq 0$. Both components are uniformly damped by the factor $(1 - \hat{\phi}_t)$, which is small when persistence is perceived to be high. Hence a transparent sufficient condition for $\partial p_t^{\text{in}*} / \partial \hat{\phi}_t < 0$ is $(1 - \hat{\phi}_t) \gamma \{ \mathcal{M}_{t-1}(\hat{\phi}_t) + \mathcal{B}_{t-1}(\hat{\phi}_t) \} < A_t$, i.e., amplification is dominated by the (legal + reputational) marginal deterrence level. In the pre-2011 Singapore benchmark with high entering norms (so A_t large), t moderately large, and favorable recent signals (so $\mathcal{M}_{t-1} > 0$ and $\mathcal{B}_{t-1} \geq 0$ but both scaled by small $(1 - \hat{\phi}_t)$), the inequality holds and the interior policy is decreasing in $\hat{\phi}_t$. Projecting back to $[0, 1]$ leaves these signs unchanged on the interior, while on the bounds the slope is flat until the policy re-enters the interior.

- (3) *Concluding sign guidance.* In our setting, the sign of the correlation with the persistence belief $\hat{\phi}_t$ is tightly disciplined by the following conditions: (i) For the *minimum enforcement* $\bar{p}_t = \mathbb{E}_t[B / (F + \gamma \hat{R}_{t-1})]$, the dependence runs through $\hat{R}_{t-1}(\hat{\phi}_t)$ only. Along histories with strong entering norms (i.e., $\hat{R}_{t-1}$ high at the benchmark) and early transition (t small), $\hat{R}_{t-1}$ increases in $\hat{\phi}_t$, so $\bar{p}_t$ decreases in $\hat{\phi}_t$. (ii) For the *interior policy* $p_t^{\text{in}*} = \frac{\rho}{\kappa} \mathbb{E}_t[(1 - \hat{\phi}_t)(F + \gamma \hat{R}_{t-1})]$, the total slope obeys

$$\frac{\partial p_t^{\text{in}*}}{\partial \hat{\phi}_t} = \frac{\rho}{\kappa} \mathbb{E}_t \left[\underbrace{-(F + \gamma \hat{R}_{t-1})}_{\text{substitution}} + \underbrace{(1 - \hat{\phi}_t) \gamma \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t}}_{\text{reputational amplification}} \right].$$

Hence $p_t^{\text{in}*}$ is *decreasing* in $\hat{\phi}_t$ whenever

$$(1 - \hat{\phi}_t) \gamma \frac{\partial \hat{R}_{t-1}}{\partial \hat{\phi}_t} < F + \gamma \hat{R}_{t-1},$$

which holds generically (i) as $\hat{\phi}_t \rightarrow 1$ (the amplification vanishes), and (ii) when entering norms are high so that A_t is large, and the memory-weighted drift of reputation is moderate (so that

$(1 - \hat{\phi}_t)\gamma \partial \hat{R}_{t-1} / \partial \hat{\phi}_t < A_t$), both $\bar{p}_t$ and $p_t^{\text{in}*}$ decrease in $\hat{\phi}_t$ under our empirical narrative (high $\tilde{R}_0$). Overall, under our empirical narrative (Singapore in 2000s with presumably high norm and persistence), both the floor $\bar{p}_t$ and the interior policy $p_t^{\text{in}*}$ move *down* with higher perceived persistence.

Remarks. Which mechanism dominates—substitution or reputational amplification—is ultimately an empirical question. First, Singapore plausibly operated near the honesty corner in the late 2000s ($x \approx 1$), where small adverse movements in the reputational signal can trigger policy adjustments. Second, as discussed in Section III.A.2 and formalized in Test 2, the 2011 tightening followed adverse innovations ($\tilde{R}_t < \mathbb{E}_{t-1}[\tilde{R}_t]$) that rationally shift beliefs about persistence downward for moderate noise (i.e., $\Delta \hat{\phi}_t < 0$). Near the corner, reputational amplification is attenuated by the $(1 - \hat{\phi}_t)$ factor and small behavioral responses, while substitution is large when $A_t = F + \gamma \hat{R}_{t-1}$ is high. The observed response (bad signal $\Rightarrow$ stronger enforcement) is therefore more naturally explained by substitution dominating amplification. Guided by this logic, in what follows we adopt the maintained hypothesis $\partial p_t^* / \partial \hat{\phi}_t < 0$ in the baseline and assess robustness where amplification could be strong.

E2. Analysis on Question 1

Fix a period t and treat $(p_t, \hat{R}_{t-1})$ as given. Ignoring the small ϵ_t , the rank- i honesty share is

$$x_{i,t} = \left[1 - B_i + p_t(F_i + \gamma_i \hat{R}_{t-1}) \right]_{[0,1]} = \left[1 - B_i + D_{i,t} \right]_{[0,1]}, \quad D_{i,t} \equiv p_t(F_i + \gamma_i \hat{R}_{t-1}).$$

Equivalently, on the interior ($0 < x_{i,t} < 1$),

$$x_{i,t} = 1 - B_i + D_{i,t}, \quad \frac{\partial x_{i,t}}{\partial p_t} = F_i + \gamma_i \hat{R}_{t-1} > 0, \quad \frac{\partial^2 x_{i,t}}{\partial p_t \partial \hat{R}_{t-1}} = \gamma_i > 0.$$

Corners obtain when the bracketed term exits $[0, 1]$:

$$\underbrace{B_i \leq D_{i,t}}_{\text{full compliance}} \Rightarrow x_{i,t} = 1, \quad \underbrace{B_i \geq 1 + D_{i,t}}_{\text{full misconduct}} \Rightarrow x_{i,t} = 0.$$

Define the *rank-specific compliance threshold*

$$\bar{p}_{i,t} \equiv \frac{B_i}{F_i + \gamma_i \hat{R}_{t-1}} \quad (\text{minimum } p_t \text{ to attain } x_{i,t} = 1).$$

Since policy is uniform, heterogeneity across ranks operates via (B_i, F_i, γ_i) .

Misconduct mass and where it concentrates. Let $c_{i,t} \equiv 1 - x_{i,t}$ denote the corruption share. On the interior, $c_{i,t} = B_i - D_{i,t}$; at corners, $c_{i,t} \in \{0, 1\}$ as above. Thus, for given $(p_t, \hat{R}_{t-1})$,

$$c_{i,t} = \begin{cases} 1, & B_i \geq 1 + D_{i,t} \quad (\text{fully corrupt}); \\ B_i - D_{i,t}, & D_{i,t} < B_i < 1 + D_{i,t} \quad (\text{interior}); \\ 0, & B_i \leq D_{i,t} \quad (\text{fully compliant}). \end{cases}$$

Aggregate misconduct is $c_t = \sum_i \alpha_i c_{i,t}$ and its cross-rank allocation is pinned down by the ordering of the *return–deterrence gaps* $B_i - D_{i,t}$.

Two benchmark regimes (Question 1).

(1) *Extreme-concentration regime (misconduct at the tails).* Suppose for some $(p_t, \hat{R}_{t-1})$,

$$B_L - D_{L,t} \gg 0, \quad B_M - D_{M,t} \leq 0, \quad B_H - D_{H,t} \gg 0.$$

This arises, for instance, when:

- Low rank has weak deterrence (F_L and γ_L small) so $D_{L,t}$ is low even under uniform p_t ;
- Middle rank faces moderate benefits B_M and nontrivial stakes $F_M + \gamma_M \hat{R}_{t-1}$, yielding $D_{M,t} \gtrsim B_M$;
- High rank enjoys large private returns B_H that outweigh even large reputational stakes $\gamma_H \hat{R}_{t-1}$, so $B_H - D_{H,t} > 0$.

Then $x_{M,t} \approx 1$ while $x_{L,t}, x_{H,t} < 1$, so misconduct concentrates at the extremes.

(2) *Middle-concentration regime (misconduct in the middle).* Suppose instead

$$B_L \leq D_{L,t}, \quad B_M > D_{M,t}, \quad B_H \leq D_{H,t}.$$

This obtains when:

- Low rank has small illicit returns B_L (e.g., limited scope/opportunity, or high shadow cost) so compliance is easy;
- High rank faces large sanctions and/or reputational stakes (F_H and γ_H high) so $D_{H,t}$ is large;
- Middle rank sits at an interior gap $B_M - D_{M,t} > 0$.

Then $x_{L,t} = x_{H,t} \approx 1$ while $x_{M,t} < 1$, so misconduct concentrates in the middle.

Comparative statics and empirical mapping. Because $D_{i,t}$ rises in both p_t and $\hat{R}_{t-1}$, higher enforcement and stronger norms jointly push ranks toward full compliance, with larger marginal effects where γ_i is higher. The cross-rank pattern thus hinges on the ordering of

$$z_{i,t} \equiv \frac{B_i}{F_i + \gamma_i \hat{R}_{t-1}}.$$

For any given p_t , ranks with larger $z_{i,t}$ are farther from compliance ($x_{i,t}$ lower). Hence:

- If $z_{L,t}$ and $z_{H,t}$ exceed $z_{M,t}$ (e.g., weak low-rank stakes; huge high-rank benefits), misconduct concentrates at the extremes.
- If $z_{M,t}$ exceeds both $z_{L,t}$ and $z_{H,t}$ (e.g., low returns at the bottom; strong reputational/sanction stakes at the top), misconduct concentrates in the middle.

These static inequalities exactly underlie Test 1: the observed locus of misconduct across ranks reveals the relative ordering of B_i versus $F_i + \gamma_i \hat{R}_{t-1}$, and therefore where policy should target when enforcement is uniform and information on rank-specific primitives is unavailable.

E3. Analysis on Question 2

Government reactions following a negative surprise. We fix the true persistence $\phi \in (0, 1)$ (unknown to the government) and study how a *negative surprise* at the end of period t , $\Delta_t \equiv \tilde{R}_t - \mathbb{E}_t[\tilde{R}_t] < 0$, affects enforcement and honesty through belief revisions about ϕ . We compare two priors: (i) a *correct prior* ($\mathbb{E}_t[\hat{\phi}_t] = \phi$) that will be revised downward and thus *underestimates* persistence after the surprise; and (ii) an *over-optimistic prior* ($\mathbb{E}_t[\hat{\phi}_t] > \phi$) that will be revised downward *toward* the truth. We first focus on general cases with interior solutions, and assume ϵ_t is small. To fit our analysis with the case of Singapore around late 2000s, we assume that pre-period reputation is high, and regularity conditions discussed in Appendix E1 are satisfied, so that p_t^* decreases in $\hat{\phi}_t$.

Period t : policy and honesty under alternative priors. Recall that the optimal enforcement at t is

$$p_t^* = \min \left\{ p_t^{\text{in}*} = \mathbb{E}_t \left[\frac{(1 - \hat{\phi}_t) \rho(F + \gamma \hat{R}_{t-1}(\hat{\phi}_t))}{\kappa} \right]_{[0,1]}, \bar{p}_t = \mathbb{E}_t \left[\frac{B}{F + \gamma \hat{R}_{t-1}(\hat{\phi}_t)} \right] \right\}.$$

Given $p_t^{\text{in}*}$, the (interior) honesty share is

$$x_t = 1 - B + p_t^{\text{in}*} (F + \gamma \mathbb{E}_t[\hat{R}_{t-1}]) + \epsilon_t.$$

Hence:

$$(\text{Correct prior}) \quad \hat{\phi}_t = \phi \Rightarrow x_t^C = 1 - B + \mathbb{E}_t \left[\frac{(1 - \phi) \rho(F + \gamma \hat{R}_{t-1}(\phi))}{\kappa} \right] (F + \gamma \mathbb{E}_t[\hat{R}_{t-1}(\phi)]) + \epsilon_t;$$

$$(\text{Over-optimistic prior}) \quad \hat{\phi}_t = \phi + \delta, \delta > 0$$

$$\Rightarrow x_t^O = 1 - B + \mathbb{E}_t \left[\frac{(1 - \phi - \delta) \rho(F + \gamma \hat{R}_{t-1}(\phi + \delta))}{\kappa} \right] (F + \gamma \mathbb{E}_t[\hat{R}_{t-1}(\phi + \delta)]) + \epsilon_t.$$

As discussed above, the substitution effect dominates if the pre-period norm level is high enough. Therefore $x_t^O < x_t^C$ (pre-surprise corruption is higher under over-optimism).

Belief revision after the negative surprise and period $t + 1$ policy. The surprise $\Delta_t < 0$ leads to a downward revision $\hat{\phi}_{t+1} < \hat{\phi}_t$. Denote the posterior reputation estimate by $\hat{R}_t$ (lower than its prior

forecast due to $\Delta_t < 0$). The period- $t + 1$ interior policy is

$$p_{t+1}^{\text{in}*} = \mathbb{E}_{t+1} \left[\frac{(1 - \hat{\phi}_{t+1}) \rho(F + \gamma \hat{R}_t)}{\kappa} \right],$$

and the honesty share is

$$x_{t+1} = 1 - B + p_{t+1}^{\text{in}*} (F + \gamma \mathbb{E}_{t+1}[\hat{R}_t]) + \epsilon_{t+1} = 1 - B + \mathbb{E}_{t+1} \left[\frac{(1 - \hat{\phi}_{t+1}) \rho(F + \gamma \hat{R}_t)}{\kappa} \right] (F + \gamma \mathbb{E}_{t+1}[\hat{R}_t]) + \epsilon_{t+1}.$$

Comparative statics at $t + 1$ (interior) and sufficient conditions. Write

$$x_{t+1} = 1 - B + \frac{\rho}{\kappa} (F + \gamma \mathbb{E}_{t+1}[\hat{R}_t]) \left[(F + \gamma \mathbb{E}_{t+1}[\hat{R}_t]) - \mathbb{E}_{t+1}[\hat{\phi}_{t+1} (F + \gamma \hat{R}_t)] \right] + \epsilon_{t+1}.$$

Decompose the posterior cross-moment as

$$\mathbb{E}_{t+1}[\hat{\phi}_{t+1} (F + \gamma \hat{R}_t)] = F \mathbb{E}_{t+1}[\hat{\phi}_{t+1}] + \gamma \mathbb{E}_{t+1}[\hat{\phi}_{t+1}] \mathbb{E}_{t+1}[\hat{R}_t] + \gamma \text{Cov}_{t+1}(\hat{\phi}_{t+1}, \hat{R}_t).$$

To obtain transparent signs, it is convenient to temporarily ignore $\text{Cov}_{t+1}(\hat{\phi}_{t+1}, \hat{R}_t)$. Then a first-order change around $(\hat{\phi}, \hat{R})$ gives

$$\frac{\partial x_{t+1}}{\partial \mathbb{E}_{t+1}[\hat{\phi}_{t+1}]} \approx -\frac{\rho}{\kappa} (F + \gamma \mathbb{E}_{t+1}[\hat{R}_t])^2, \quad \frac{\partial x_{t+1}}{\partial \mathbb{E}_{t+1}[\hat{R}_t]} \approx 2\frac{\rho}{\kappa} \gamma (F + \gamma \mathbb{E}_{t+1}[\hat{R}_t]) (1 - \mathbb{E}_{t+1}[\hat{\phi}_{t+1}]).$$

Hence, following a negative surprise $\Delta_t < 0$ that induces $\Delta \mathbb{E}_{t+1}[\hat{\phi}_{t+1}] < 0$ and $\Delta \mathbb{E}_{t+1}[\hat{R}_t] < 0$, honesty at $t + 1$ rises ($\Delta x_{t+1} > 0$) if and only if the *substitution* gain from the downward revision in persistence dominates the *reputational* drag from the lower $\hat{R}_t$:

$$(9) \quad |\Delta \mathbb{E}_{t+1}[\hat{\phi}_{t+1}]| > \frac{\gamma(2 - \mathbb{E}_{t+1}[\hat{\phi}_{t+1}])}{F + \gamma \hat{R}_t} |\Delta \hat{R}_t|.$$

We note that allowing for $\text{Cov}_{t+1}(\hat{\phi}_{t+1}, \hat{R}_t)$ makes this condition harder to satisfy since $\text{Cov}_{t+1}(\hat{\phi}_{t+1}, \hat{R}_t)$ tends to be positive.

Case (i): Correct prior revised downward. With $\hat{\phi}_t = \phi$ and $\hat{\phi}_{t+1} = \phi - \underline{\delta}$ ($\underline{\delta} > 0$), the posteriors satisfy

$$x_{t+1}^C = 1 - B + \frac{\rho}{\kappa} (F + \gamma \hat{R}_t^C) [(F + \gamma \hat{R}_t^C) - (\phi - \underline{\delta})F - \gamma \mathbb{E}_{t+1}[\hat{\phi}_{t+1} \hat{R}_t]] + \epsilon_{t+1},$$

with $\hat{R}_t^C < \mathbb{E}_t[R_t]$ because $\Delta_t < 0$. For high pre-period norm level and modest reputation shortfall $|\Delta \hat{R}_t^C|$, the condition above typically holds and $x_{t+1}^C \gtrsim x_t^C$ (a mild, transitory tightening).

Case (ii): Over-optimistic prior revised toward the truth. With $\hat{\phi}_t = \phi + \bar{\delta}_t$ and $\hat{\phi}_{t+1} = \phi + \bar{\delta}_{t+1}$ where $0 \leq \bar{\delta}_{t+1} < \bar{\delta}_t$, we have

$$x_{t+1}^O = 1 - B + \frac{\rho}{\kappa} (F + \gamma \hat{R}_t^O) [(F + \gamma \hat{R}_t^O) - (\phi + \bar{\delta}_{t+1})F - \gamma \mathbb{E}_{t+1}[\hat{\phi}_{t+1} \hat{R}_t]] + \epsilon_{t+1},$$

with $\hat{R}_t^O < \hat{R}_t^C$ because $x_t^O < x_t^C$. Here both forces are stronger: the substitution term increases since $|\Delta \mathbb{E}_{t+1}[\hat{\phi}_{t+1}]| = \bar{\delta}_t - \bar{\delta}_{t+1}$ is larger, while the reputational drag also increases because $\hat{R}_t^O$ is lower. For high pre-period norm level (so $F + \gamma \hat{R}_t$ is large), the left-hand side of (9) dominates, implying $x_{t+1}^O - x_t^O > 0$ and, quantitatively, a larger rise than in Case (i).

Within-scenario implications.

- Correct prior $\rightarrow$ underestimation. The drop in $\hat{\phi}$ is purely surprise-driven (no misspecification to correct). For high pre-period norm level, the reputation shortfall $|\Delta \hat{R}_t|$ is modest, so we expect $x_{t+1}^C \gtrsim x_t^C$ even as the government becomes temporarily too pessimistic.
- Over-optimism $\rightarrow$ correction toward ϕ . Here $|\Delta \hat{\phi}_{t+1}| = \bar{\delta}_t - \bar{\delta}_{t+1}$ is typically *larger* than in the correct-prior case, strengthening substitution; yet $\hat{R}_t^O$ is also *lower* (because $x_t^O < x_t^C$), weakening reputation. For high pre-period norm level, $(F + \gamma \hat{R}_t)$ is large, making $\Delta x_{t+1}^O > 0$ the empirically relevant prediction: x_{t+1}^O rises sharply as the government corrects under-enforcement.

We then discuss a simpler case of corner solution. Assume that the optimal enforcement under correct belief of ϕ is precisely $\bar{p}_t = \mathbb{E}_t \left[\frac{B}{F + \gamma \hat{R}_{t-1}(\phi)} \right]$. Then the over-optimism gives $x_t^O < x_t^C \approx 1$, where the correct prior gives near-zero corruption. Next, an over-reaction keeps the near-zero corruption for period $t + 1$, while the correction toward ϕ reduces, and may eliminate, the corruption in $t + 1$. Therefore, this special case yields similar implications as the general one.

Cross-scenario comparisons (Question 2).

- (1) *Pre-tightening level at t* : $x_t^O < x_t^C$ —the over-optimistic government under-enforces at t , yielding elevated misconduct before the surprise.
- (2) *Impact at $t + 1$* : Interior tightening is stronger under over-optimism (larger $|\Delta \hat{\phi}_{t+1}|$), so $x_{t+1}^O - x_t^O$ exceeds $x_{t+1}^C - x_t^C$. Persistence of the tightening further distinguishes learning from noise.

Therefore, observing (i) low and stable pre-surprise misconduct together with a transitory policy response aligns with *noise-driven overreaction* (correct prior revised too far); whereas (ii) elevated pre-surprise misconduct and a sustained post-surprise tightening with $\Delta x_{t+1} > 0$ are diagnostic of *correct learning* from an over-optimistic prior—exactly the distinction in *Question 2*.