

**From Oil to Health: Evaluating the Socio-Economic Impact of Vision 2030
on the Saudi Healthcare System**

\

Master of Healthcare Administration

Prepared by the student:

Abrar Awad Omar Al-Nakhbi

Supervision:

Dr. Al-Faisal Abdulhamid

Dr. Fayez Jarrad

Academic Year: 2026

Abstract

Saudi Arabia's Vision 2030 has pursued a decade-long transformation of its healthcare system, moving from a state-funded model toward one built on corporatized delivery through the Health Holding Company and regional Health Clusters, financed through mandatory cooperative health insurance. This thesis evaluates that reform against its own stated socio-economic goals: reduced government health expenditure, improved patient satisfaction, narrower urban-rural disparities, and lower household out-of-pocket costs. Because primary survey and interview data fell outside this project's scope, the thesis uses a desk-based, secondary-data design, drawing on World Bank, GASTAT, Ministry of Health, and Council of Health Insurance statistics alongside the published literature on Saudi and Gulf healthcare reform, and tests four hypotheses derived from the reform's own claims.

The results do not support a clean confirmation of any of the four. Government health expenditure has not declined since the Health Sector Transformation Program's 2021 launch, whether measured as a share of GDP or per capita, and this finding is read against HSTP's own five-year implementation horizon (2021–2026), which is now largely elapsed, rather than Vision 2030's more distant 2030 target year. A robustness check applying a real-terms, CPI-based deflation and a narrower post-reform window confirms the same finding under stricter terms, though the size of the apparent rise is smaller once inflation and the two years before the Health Holding Company existed are removed. Out-of-pocket spending is method-sensitive: a simple endpoint comparison shows spending back near its 2016 level, while a period-average comparison, applying the same method used for the efficiency hypothesis, shows an approximately 12% relative decline; this thesis reports both rather than selecting one. The evidence on patient satisfaction is too fragmented across independent studies to support a system-wide verdict, though implementation-focused literature identifies workforce

shortages and staff resistance to change as a plausible partial explanation. The equity hypothesis, as originally framed around geographic infrastructure exclusion, is not well supported; the literature instead points to a digital literacy and trust gap that cuts across geography rather than a purely rural-urban one, alongside a narrower resourcing gap in remote regions distinct from network connectivity. A related equity gap along citizen/non-citizen lines, visible in insurance enrollment volatility among expatriate residents, was identified but not formally analyzed; this thesis also finds that an expatriate dependent levy phased in from 2017, rather than the COVID-19 pandemic alone, offers a better-timed explanation for part of that volatility. The thesis also identifies a material discrepancy between the life expectancy baseline Vision 2030 cites and the baseline shown in international data for the same year, and acknowledges, as a limitation, that it has not applied a fully consistent standard of skepticism to Vision 2030's own self-reported statistics throughout.

These findings suggest the reform's institutional apparatus, the Health Holding Company, the Health Clusters, and the expanded insurance network, is in place, but the outcomes it was designed to produce have not yet clearly and consistently appeared in the public record.

Keywords: Vision 2030; healthcare reform; Saudi Arabia; health sector transformation; rentier state; healthcare privatization; out-of-pocket expenditure; health equity

Data Availability Statement

All data used in this thesis are drawn from publicly accessible sources: the World Bank World Development Indicators, GASTAT publications, Saudi Ministry of Health statistical releases, and Council of Health Insurance annual reports, alongside the published academic literature cited throughout. No primary data were collected. Full source details, access points, and retrieval information are documented in Appendix A, and the complete year-by-year indicator series underlying every table and figure in Chapter 4 is consolidated in Appendix B.

One data point (Council of Health Insurance enrollment figures, Chapter 4, Table 4.2) is compiled from secondary industry reporting rather than a single primary publication; this limitation is discussed in Chapter 5, Section 5.4 and Appendix A.

Table of Contents

Acknowledgments.....	Error! Bookmark not defined.
Abstract.....	2
Data Availability Statement.....	3
Table of Contents	4
List of Tables	6
List of Figures	7
List of Abbreviations.....	8
Chapter 1: Introduction	8
1.1 Background and Context	8
1.2 Problem Statement	11
1.3 Significance of the Study.....	13
1.4 Objectives of the Study.....	14
1.5 Research Questions	15
1.6 Research Hypotheses	15
1.7 Scope and Limitations	16
1.8 Structure of the Thesis.....	17
Chapter 2: Literature Review and Theoretical Framework.....	19
2.1 Rentier State Theory and the Post-Rentier Transition.....	19

2.2 Comparative Perspectives on Healthcare Privatization.....	21
2.3 The Saudi Journey: From Public Provision to Corporatized Care	24
2.4 Health Insurance, Corporatization, and Out-of-Pocket Spending.....	26
2.5 The Digital Frontier: Telehealth, Sehaty, and the Digital Divide	28
2.6 Theoretical Framework.....	29
2.6.1 Theory of Change	29
2.6.2 Agency Theory and Market Failure in Health Economics	30
2.6.3 Social Determinants of Health and Equity.....	30
2.7 Synthesis and Gap Statement.....	32
Chapter 3: Methodology.....	35
3.1 Research Philosophy and Design	35
3.2 Data Sources	36
3.3 Document and Data Inclusion Criteria	37
3.4 Analytical Approach by Hypothesis.....	38
3.5 Validity and Reliability.....	41
3.6 Limitations of the Design	43
Chapter 4: Results and Analysis	45
4.1 Hypothesis 1: Efficiency.....	45
Robustness Check: Real-Terms and Window-Sensitivity Analysis	50
4.2 Hypothesis 2: Quality and Patient Satisfaction.....	53
4.3 Hypothesis 3: Equity (Urban–Rural Gap)	55

4.4 Hypothesis 4: Insurance and Out-of-Pocket Spending.....	57
4.5 The Life Expectancy Baseline	63
4.6 Summary of Verdicts.....	66
Chapter 5: Discussion and Conclusion.....	69
5.1 Summary of Findings	69
5.2 Interpretation Through the Theoretical Framework.....	70
5.3 Contribution and Positioning	72
5.4 Limitations	74
5.5 Robustness and Alternative Explanations.....	75
5.6 Policy Recommendations	76
5.7 Directions for Future Research	79
5.8 Conclusion	81
References	82
Appendix A: Data Retrieval Details	87
Appendix B: Full Underlying Data Series	90

List of Tables

Table 1.1: Vision 2030 Healthcare Reform Timeline (Chapter 1)

Table 2.1: Comparative GCC Healthcare Financing Models (Chapter 2)

Table 2.2: Theoretical Framework Application Map (Chapter 2)

Table 3.1: Secondary data sources (Chapter 3)

Table 3.2: Data Availability and Evidentiary Strength by Hypothesis (Chapter 3)

Table 4.1: Saudi Arabia government and total health expenditure, 2016–2023 (Chapter 4)

Table 4.2: Out-of-pocket health expenditure and insurance enrollment, 2016–2024 (Chapter 4)

Table 4.3: Hypothesis Findings Summary (Chapter 4)

Table 4.4: Real-terms and window-sensitivity comparison (Chapter 4)

Table 5.1: Policy Recommendations (Chapter 5)

Table B.1: Consolidated indicator series, Saudi Arabia, 2016–2024 (Appendix B)

List of Figures

Figure 2.1: Literature Base by Publication Year (Chapter 2)

Figure 4.1: Saudi Arabia Health Expenditure as % of GDP, 2016–2023 (Chapter 4)

Figure 4.2: Saudi Arabia Health Expenditure per Capita, 2016–2023 (Chapter 4)

Figure 4.3: Out-of-Pocket Health Expenditure, 2016–2023 (Chapter 4)

Figure 4.4: CCHI Insurance Enrollment, 2016–2024 (Chapter 4)

Figure 4.5: Government Health Expenditure (% of GDP), Saudi Arabia vs. GCC Peers, 2016–2023 (Chapter 4)

Figure 4.6: Public-Private Partnership Pipeline Progress Against Stated 2030 Target (Chapter 4)

Figure 4.7: Out-of-Pocket Health Expenditure, Saudi Arabia vs. GCC Peers, 2016–2023 (Chapter 4)

Figure 4.8: Vision 2030's Stated Life Expectancy Trajectory vs. World Bank/UN Data (Chapter 4)

Figure 4.9: Nominal vs. Real (2016-Basis) Per-Capita Health Expenditure, 2016–2023 (Chapter 4)

List of Abbreviations

Abbreviation	Meaning
CCHI	Council of Cooperative Health Insurance
CHI	Council of Health Insurance
GASTAT	General Authority for Statistics (Saudi Arabia)
GCC	Gulf Cooperation Council
GDP	Gross Domestic Product
HHC	Health Holding Company
HSTP	Health Sector Transformation Program
KPI	Key Performance Indicator
MENA	Middle East and North Africa
MOH	Ministry of Health (Saudi Arabia)
NPM	New Public Management
OOP	Out-of-Pocket (expenditure)
PPP	Public-Private Partnership
SDG	Sustainable Development Goal
SDOH	Social Determinants of Health
UHC	Universal Health Coverage
WHO	World Health Organization

Chapter 1: Introduction

1.1 Background and Context

Saudi Arabia's modern state emerged on a fiscal foundation that political economists call the rentier model: a government funded primarily by external revenue (oil), rather than domestic taxation, and which in turn distributes that revenue to citizens through free or heavily subsidized public services (Beblawi & Luciani, 1987). Healthcare was one of the clearest expressions of that arrangement. From the 1970s onward, the Ministry of Health (MOH) built out a national network of hospitals and primary care centers, funded them from the state budget, staffed them largely with expatriate physicians, and offered care to citizens at no direct cost. The results were real: infant mortality fell sharply, and the infectious diseases that had once been endemic to the Peninsula were brought under control (Walston, Al-Harbi, & Al-Omar, 2008).

That model has run into two structural limits. The first is fiscal: oil price volatility and the global shift toward renewable energy make a health budget financed almost entirely by

hydrocarbon rents an increasingly unstable proposition. The second is epidemiological. As the population has grown and life expectancy has risen, the disease burden has shifted from infectious to chronic and lifestyle-related conditions such as diabetes, cardiovascular disease, and obesity, which are expensive to manage over decades rather than treat once (Walston et al., 2008). A system built to expand access rapidly was not built to manage this kind of long-run cost curve.

Vision 2030, launched in April 2016, is the Saudi government's response to both pressures at once. Its stated aim is to diversify the economy away from oil dependency and toward a private-sector-driven model, and healthcare sits inside that broader strategy in a dual role: as a social service that must remain accessible, and as a sector the government wants to convert into a source of private investment, employment, and non-oil GDP growth (Vision 2030, 2016). The operational vehicle for this shift in health is the Health Sector Transformation Program (HSTP), one of the Vision Realization Programs, first launched in 2021 with a five-year horizon, i.e., to 2026 (Ministry of Health, 2024a; see Table 1.1 for the full timeline). Under HSTP, the MOH's role is meant to narrow from funder-provider-regulator to regulator alone. Day-to-day operation of public hospitals and clinics has been transferred to a new entity, the Health Holding Company (HHC), established by royal decree in 2022, which in turn oversees regional Health Clusters structured as autonomous regional operating entities rather than line departments of a ministry, described in the literature as resembling accountable-care-style organizations though without the specific shared-savings and attribution mechanics of the US Medicare model that term originally denotes (Chowdhury, Mok, & Leenen, 2021). Financing is being reorganized around mandatory cooperative health insurance, administered by the Council of Health Insurance (CHI, previously known as the Council of Cooperative Health Insurance, or CCHI, the name still used in earlier data series and by data compilers tracking the scheme under its original name; see Chapter 4, Table 4.2),

which describes its own regulatory role and market-structure data in its most recent annual report (Council of Health Insurance, 2023), and around public-private partnerships intended to bring in capital the state budget would otherwise have to supply directly.

The "Vibrant Society" pillar of Vision 2030 sets a concrete target for this transformation: raising average life expectancy from 74 to 80 years. That target is not pursued through hospital construction alone. It rests partly on digital tools, namely the Sehaty and Nafis platforms, which give citizens direct access to appointments, records, and teleconsultation, and partly on a broader push toward preventive care and individual responsibility for health outcomes, a departure from a system built around curative, in-person treatment.

Table 1.1: Vision 2030 Healthcare Reform Timeline

Year	Milestone
2016	Vision 2030 launched (Vision 2030, 2016)
2017	Monthly dependent levy on expatriate workers' families begins phasing in, SR100 per dependent rising to SR400 by 2020 (Arab News, 2016; see Chapter 4, Section 4.4)
2018	Value-added tax (VAT) introduced at 5%
2020	VAT tripled to 15%; COVID-19 pandemic disrupts healthcare utilization worldwide
2021	Health Sector Transformation Program (HSTP) launched with a stated five-year horizon (i.e., to 2026)
2022	Health Holding Company (HHC) established by royal decree; regional Health Clusters begin operating under HHC oversight
2023	HSTP Annual Report published (Vision 2030 Delivery Unit, 2023); Council of Ministers Decision No. 85/2023 transfers insurance-sector oversight functions from CHI/SAMA toward the newly established Insurance Authority
2024	MOH Statistical Yearbook 2024 and GASTAT Health Determinants Statistics 2024 published
2026	HSTP's own stated five-year implementation horizon concludes
2030	Vision 2030's overall target year (life expectancy target, private-sector delivery share target, PPP investment target)

This timeline matters for how "on schedule" claims are read later in the thesis (Chapters 4 and 5): HSTP itself, the specific operational program this thesis evaluates, is scoped to 2021–2026, a materially shorter and nearer-term horizon than Vision 2030's overall 2030 endpoint.

Findings about HSTP's outcomes are considered against both horizons explicitly, rather than treating the two as interchangeable.

Roughly a decade into this transformation, the volume of description of what Vision 2030's health reforms are supposed to do considerably exceeds the volume of evidence on what they have actually done. That gap is the starting point for this thesis.

1.2 Problem Statement

The reform now underway asks Saudi citizens to relate to their health system differently than they have for two generations. Care is now a service mediated by mandatory insurance and a regulated private market, delivered by a state that has stepped back from direct provision, in place of the unconditional entitlement it was before. That is a substantial change to what has functioned as an informal social contract, and it raises questions that are economic, social, and administrative in roughly equal measure.

On the economic side, the central claim behind HSTP is that corporatizing hospital operations and shifting financing toward insurance will reduce the fiscal burden on the state while improving efficiency. Whether that claim holds is, at this point, an empirical question rather than a settled fact. A corporatized system could equally shift costs sideways, onto private insurers, onto employers who fund employee coverage, and eventually onto out-of-pocket household spending, without reducing overall system cost at all. There is also a standard risk in this kind of transition: if the new regional Health Clusters or private entrants gain effective monopoly power over specific services or geographies, prices can rise rather than fall, producing a system that is more expensive to access rather than more efficient.

On the social side, a rapid corporatization of a public good tends to produce uneven results across geography and income unless it is actively managed against that outcome. Riyadh, Jeddah, and Dammam are already visibly ahead in private hospital investment and in the rollout of new Health Cluster infrastructure; whether the same holds for the Northern

Borders, Asir, or other less urbanized regions is a separate and open question. Sehaty and Nafis illustrate a parallel risk in digital form: tools that meaningfully improve access for a digitally fluent, urban population can function as a barrier for citizens who are older, less connected, or unfamiliar with smartphone-based services, if the digital channel becomes the primary channel rather than one option among several.

On the administrative side, the question is whether the Health Holding Company and its regional clusters have gained real operating autonomy or only a formal legal status that leaves substantive control with central government, the same distinction Preker and Harding's (2003) comparative work on hospital corporatization identifies as the point where reforms of this type most often fall short of their own design (Chapter 2, Section 2.2). A cluster with formal autonomy but no practical latitude over staffing, procurement, or budget allocation would behave, in the outcomes it produces, much like the ministry department it formally replaced, even though the organizational chart has changed. This is a harder question to answer from public data than the economic or social dimensions above, since it turns on internal governance practices that are not routinely disclosed, and it is one this thesis's desk-based design cannot resolve directly; the efficiency and quality findings in Chapter 4 are read against this uncertainty rather than treating the formal creation of the Health Holding Company as proof that operational autonomy exists in practice.

The research reviewed for this thesis (see Chapter 2) shows that these questions have not gone unaddressed in the academic literature: there is a growing body of published work describing HSTP's institutional design, documenting the creation of the HHC and Health Clusters, and reviewing the insurance system's coverage expansion. What is comparatively scarce is empirical synthesis that tests the reform's own stated claims (efficiency, quality, equity, and financial protection) against longitudinal data, rather than describing the policy architecture on its own terms. This thesis is positioned in that narrower gap: not as the first

study of Vision 2030's health reforms, but as an attempt to evaluate them against the economic and social outcomes the reform explicitly promised to deliver.

1.3 Significance of the Study

Theoretical significance: most of the scholarship on healthcare privatization and corporatization is built on Western, already-industrialized cases: the UK's NHS reforms, Australia's system, and the wave of restructuring across Latin America and Eastern Europe in the 1990s. Saudi Arabia is a different case: a resource-rich state voluntarily dismantling elements of its own welfare provision while still possessing the fiscal capacity to fund the old model, motivated by long-run sustainability rather than short-run crisis. This case sits within a growing, though still developing, literature on "late rentierism" (Gray, 2011) and Saudi-specific reform political economy (Hertog, 2024), which this thesis draws on rather than treats as absent; the more precise claim, developed further in Chapter 2, is that this literature is well populated on institutional description but thin specifically on empirical, outcome-focused evaluation of the health sector, which is the gap this thesis addresses. Consistent with the methodological commitment in Chapter 3 not to generalize findings outward without a country-specific anchor, this thesis's claims are scoped to Saudi Arabia; any relevance to other Gulf states pursuing similar diversification agendas is left for future comparative work rather than asserted here. The one partial exception, discussed in Chapter 2, Section 2.2, is the regional health-economics-capacity literature (AlRuthia et al., 2025), which situates Saudi Arabia's specific PPP pipeline shortfall (Chapter 4, Section 4.1) inside a documented GCC-wide pattern rather than treating it as a country-specific anomaly; even there, this thesis draws on the regional literature to interpret a Saudi-specific finding, rather than extending the Saudi finding outward to claim anything about the other five GCC states.

Practical significance: for the Ministry of Health, the Health Holding Company, and the Council of Health Insurance, an independent, data-grounded assessment of where the reform

is and is not meeting its own targets functions as a diagnostic rather than an endorsement or an indictment: the kind of input that allows a correction mid-course rather than after the 2030 deadline has passed. For investors evaluating the Saudi healthcare market, the same analysis clarifies where demand, insurance coverage, and regulatory risk actually stand, as opposed to where the policy documents project them to stand. For Saudi citizens, whose relationship with the health system is what is actually changing, this thesis attempts to hold the reform to the concrete outcomes (cost, access, quality) that determine whether the transformation is felt as an improvement or as a cost shifted downward. For the roughly one-third of Saudi Arabia's population that is expatriate, and whose insurance coverage is tied to formal employment rather than citizenship, this thesis's finding on the dependent-levy policy's effect on CCHI enrollment (Chapter 4, Section 4.4) is relevant in its own right: it identifies a specific, dated policy lever, separate from the reform's healthcare-specific mechanisms, with a demonstrable bearing on coverage for a population without the citizen-facing MOH fallback. The analysis is framed against UN Sustainable Development Goal 3 (health and well-being) and Goal 10 (reduced inequality), both of which the Saudi government has itself invoked in connection with Vision 2030.

1.4 Objectives of the Study

This thesis pursues five objectives:

1. To assess whether the shift of operational authority from the MOH to the Health Holding Company and regional Health Clusters is associated with a measurable change in government health expenditure.
2. To evaluate how the expansion of mandatory cooperative health insurance has affected out-of-pocket healthcare spending by Saudi households.
3. To assess whether the shift toward value-based, corporatized care is associated with improvements in patient-reported quality and satisfaction, drawing on the existing body of published survey research.

4. To examine whether digital health platforms (Sehaty, Nafis) and broader Health Cluster investment are distributed evenly across urban and rural regions, or whether they track existing geographic advantage.
5. To identify the structural and data-availability constraints that limit a full assessment of HSTP's socio-economic impact at this stage, and to set out what a future, primary-data study would need to resolve them.

1.5 Research Questions

Economic

6. Is the shift toward regional Health Clusters and public-private partnerships associated with a reduction in government health expenditure as a share of GDP or national budget?
7. How has the expansion of mandatory health insurance affected household out-of-pocket healthcare spending, and what does that trend suggest about investor confidence in the sector?

Social

8. Does the existing published evidence on patient satisfaction support the claim that corporatized, value-based care has improved perceived quality of care?
9. Is there a measurable gap between healthcare access and quality in Riyadh, Jeddah, and Dammam compared with less urbanized regions?

Technological and systemic

10. Is there evidence that low digital literacy or limited connectivity constrains the reach of telehealth and digital diagnostic tools among specific population groups?
11. What structural, financial, or cultural frictions does the secondary literature and available public data identify as obstacles to the transition from a state-run to a value-based system?

1.6 Research Hypotheses

Hypothesis 1 (Efficiency).

H₀: Corporatization of healthcare delivery through Health Clusters shows no association with a reduction in government health expenditure.

H₁: Corporatization is associated with a measurable reduction in government health expenditure as a share of GDP or national budget.

Hypothesis 2 (Quality).

H₀: Published patient satisfaction data show no meaningful improvement under the corporatized, value-based care model relative to the prior MOH-run model.

H₁: Published patient satisfaction data show a measurable improvement under the corporatized model.

Hypothesis 3 (Equity).

H₀: There is no measurable difference in healthcare access, quality, or digital health service usage between urban and rural populations.

H₁: Urban populations show a measurable advantage in healthcare access, quality, and digital health service usage over rural populations.

Hypothesis 4 (Insurance).

H₀: Expansion of mandatory health insurance shows no association with a reduction in household out-of-pocket healthcare expenditure.

H₁: Expansion of mandatory health insurance is associated with a measurable reduction in household out-of-pocket healthcare expenditure.

1.7 Scope and Limitations

Delimitations: the thesis is confined to the socio-economic dimensions of Vision 2030's health reforms (cost, access, financing, and equity) and does not address clinical outcomes, treatment efficacy, or the biomedical content of care. Geographically, it covers the Kingdom as a whole, with particular attention to the urban–provincial divide. Temporally, the analysis frame runs from 2016 (the launch of Vision 2030) through the present, the period in which HSTP has moved from design to implementation; the quantitative data underlying Chapter 4 extend through 2023 for most series and through 2024 for CHI/CCHI enrollment specifically

(see Table 3.1 and Appendix A), roughly three years short of Vision 2030's own 2030 horizon, which limits how much this thesis can say about the reform's most recent phase.

Limitations: this thesis uses a desk-based, secondary-data design rather than primary survey or interview data (see Chapter 3 for the rationale and its consequences). Three limitations follow directly from that choice. First, detailed financial data from individual Health Clusters and the Health Holding Company are not publicly disclosed at a granular level; the analysis relies on aggregated national statistics and acknowledges where Health Cluster-level claims cannot be independently verified. Second, the evidence on patient satisfaction (Hypothesis 2) is drawn from studies conducted at different times, in different regions, and with different instruments; this thesis synthesizes that evidence rather than measuring satisfaction directly, and the resulting picture is necessarily less precise than a purpose-built survey would provide. Third, association observed in longitudinal public data cannot establish causation; where government health expenditure or out-of-pocket spending moves in a particular direction, this thesis reports that movement and considers competing explanations, but does not claim the reform alone produced it.

1.8 Structure of the Thesis

Chapter 2 reviews the literature on rentier-state transitions, comparative healthcare privatization, and the specific institutional history of Vision 2030's health reforms, and sets out the three-lens theoretical framework (Theory of Change, Agency Theory, and the Social Determinants of Health) used to interpret the findings. Chapter 3 details the secondary-data research design, the specific public data sources used, and their limitations. Chapter 4 presents findings organized around the four hypotheses, including a robustness check on the efficiency finding and a reframing of the insurance-enrollment discussion. Chapter 5 discusses those findings against the theoretical framework, states the study's limitations plainly, and closes with policy recommendations and directions for future primary-data

research. Two appendices follow the reference list: Appendix A documents data retrieval details, and Appendix B consolidates the full underlying data series referenced throughout Chapter 4.

Chapter 2: Literature Review and Theoretical Framework

2.1 Rentier State Theory and the Post-Rentier Transition

Beblawi and Luciani's (1987) *The Rentier State* remains the reference point for understanding how Gulf states like Saudi Arabia have historically financed public services. In their formulation, a rentier state draws the bulk of its revenue from external sources (for Saudi Arabia, oil exports) rather than from taxing domestic economic activity. Because that revenue does not depend on a productive domestic tax base, the state is able to distribute it broadly through free or subsidized public services, including healthcare, without the same political bargaining over taxation that shapes welfare provision elsewhere. The arrangement Beblawi and Luciani describe is not merely an economic structure; it functions as an implicit exchange between state and citizen, in which broad public provision substitutes for political representation.

That arrangement becomes harder to sustain as two conditions change: when the revenue base itself becomes less reliable, and when the cost of the services being distributed grows faster than the revenue funding them. Both conditions apply to Saudi Arabia's health sector at present. Oil price volatility and the global transition toward renewable energy put long-run pressure on the revenue side, while population growth, rising life expectancy, and the shift toward chronic disease put sustained upward pressure on the cost side. Gray's (2011) concept of "late rentierism" describes this specific transition: a rentier state that retains the fiscal capacity to continue distributing rents but begins restructuring how it distributes them, in anticipation of future constraint rather than in response to an immediate fiscal crisis. Hertog's (2024) analysis of Vision 2030's political economy situates the health reforms within this broader pattern, arguing that the older Saudi social contract, built on state spending and state employment, is eroding gradually rather than collapsing, which is consistent with a reform process that redesigns institutions before the fiscal pressure becomes acute rather than after.

Herb's (2014) comparative study of rentier bargains in Kuwait and the UAE, while not about Saudi Arabia or healthcare specifically, establishes a useful benchmark for this thesis: rentier states renegotiate their distributive bargains at different speeds and through different institutional channels depending on domestic political constraints, which is a caution against assuming Saudi Arabia's health-sector transition will follow the same pace or shape as reforms elsewhere in the Gulf. Vision 2030's health reforms sit at this junction, and the literature on GCC governance frames the shift underway (the move of the Ministry of Health from direct provider to regulator, with hospital operation handed to the Health Holding Company and its regional clusters) as a broader regional pattern of moving from a "provider state" toward a "regulatory state" model (Chowdhury, Mok, & Leenen, 2021).

The health-sector reforms this thesis focuses on did not arrive alone; they were part of a wider package of fiscal-consolidation measures under the same Fiscal Balance Program that introduced the VAT changes in Table 1.1, including a levy on expatriate workers' dependents phased in from 2017 (Arab News, 2016; discussed in Chapter 4, Section 4.4). None of these instruments is healthcare-specific, but together they illustrate Gray's (2011) late-rentierism pattern operating across several policy domains at once rather than in healthcare in isolation: a state broadening its non-oil revenue base and reducing the cost of servicing its non-citizen population through several distinct fiscal levers introduced within the same multi-year program, of which the healthcare reforms this thesis evaluates are one part among several. This matters for how Chapter 4's findings should be read: where a trend in the health data coincides with a milestone in Table 1.1, the coincidence is at least as likely to reflect this broader fiscal program as a healthcare-specific policy effect, and Chapter 4 treats that ambiguity explicitly rather than defaulting to a healthcare-specific explanation for every observed shift.

2.2 Comparative Perspectives on Healthcare Privatization

The theoretical vocabulary for this kind of reform predates Vision 2030 by several decades. Hood (1991) named and characterized New Public Management (NPM), the doctrine prevalent across English-speaking public administrations from the 1980s onward holding that public services perform better when structured to resemble private firms: disaggregated into semi-autonomous units, subject to performance metrics, and exposed to competitive or quasi-competitive pressure. The Health Holding Company's structure, with Health Clusters operating as autonomous regional entities rather than departments of a central ministry, is a recognizable instance of the NPM template applied to health, and more specifically an instance of what Preker and Harding's (2003) World Bank volume on hospital corporatization calls granting "corporatized" status: formal legal and managerial autonomy for public hospitals short of full privatization. That volume's central finding, drawn from comparative cases across middle-income and transition economies, is that corporatized status does not automatically translate into real operational autonomy; parent ministries and holding entities often retain effective control over budgets, staffing, and major decisions even after the formal governance structure changes. Whether the Health Holding Company and its Health Clusters have crossed that line from formal to real autonomy is not something this thesis's data can test directly, but it is the specific question Preker and Harding's framework raises and that Chapter 4's efficiency findings should be read against.

Whether that template delivers on its efficiency promise is contested rather than settled. Alayed, Alrumeh, Alkanhal, and Alhuthil's (2024) systematic review of privatization's effects on healthcare systems finds a mixed evidence base: some settings show gains in efficiency and reduced waiting times following competitive reform, while others show that privatization, absent strong regulatory capacity, produces fragmented care and rising costs rather than falling ones. The recurring risk identified across this literature is stratification: a

system in which the quality of care a patient receives correlates with income or insurance status rather than need, a pattern documented most clearly in health systems that privatized rapidly without building the regulatory infrastructure to prevent it. This is the tension the Saudi reform has to manage directly: HSTP explicitly retains the Ministry of Health as regulator specifically to avoid this outcome, but a regulator's stated intent and its actual capacity to enforce equitable access are separate questions, and the second is the one this thesis is positioned to examine.

Within the Gulf specifically, the reform pattern is not unique to Saudi Arabia. AlRuthia and colleagues' (2025) regional overview, prepared for the Gulf Health Economics Association, and government-level reporting summarized in Table 2.1 both describe a shared GCC-wide trend toward mandatory, employer-linked insurance for the expatriate workforce alongside continued free or subsidized care for citizens, with the balance of government-versus-private financing varying substantially by country. Their review frames the underlying pressure in terms directly relevant to this thesis's efficiency hypothesis: GCC states built universal, oil-funded healthcare coverage decades ago, but as population and healthcare demand keep growing, financing that coverage without weighing the value of the health technologies being reimbursed has become steadily harder to sustain. AlRuthia and colleagues identify a specific institutional gap behind that pressure: most GCC health systems, Saudi Arabia included, lack the health economics and outcomes research (HEOR) capacity to systematically evaluate whether a given technology or treatment pathway is worth its cost before it is adopted into the reimbursed benefits package, which leaves cost growth to be managed after the fact through blunter instruments such as the PPP financing push and insurance mandate this thesis evaluates directly, rather than through upstream technology assessment.

Table 2.1: Comparative GCC Healthcare Financing Models

Country	Citizen Care Model	Expatriate/Private-Sector Coverage	Government Share of Health Financing
---------	--------------------	------------------------------------	--------------------------------------

Saudi Arabia	MOH-funded free/subsidized care, transitioning toward HHC/Health Cluster delivery	Mandatory private insurance since 2023 reforms, regulated via CHI/Insurance Authority	Moderate-high (see Figure 4.5 for trend)
UAE	Free or low-cost citizen care	Mandatory in Dubai and Abu Dhabi; nationwide mandate effective January 2025	Lowest in the GCC (~52%)
Bahrain	Largely free citizen care	Expanding mandatory insurance requirements for expatriate workers	Moderate
Kuwait	Free citizen care, fixed co-payments for services	Employer-linked coverage; public-private partnership programs	Highest in the GCC (~87%)
Oman	Largely free citizen care	Expatriate insurance requirements expanding	Moderate-high
Qatar	Free citizen care, moving toward proposed universal insurance	Increased private-sector integration	Moderate

Sources: AlRuthia et al. (2025); government/regulatory reporting cited in Table 1.1 and Section 2.2 text.

Saudi Arabia's HSTP is best read as a more structurally ambitious version of a shift already underway across the GCC, rather than an isolated policy experiment. Table 2.1's government-financing-share column places Saudi Arabia in the middle of this range rather than at either extreme: Kuwait's roughly 87% government share reflects a system that has retained direct state financing more fully than its neighbors, while the UAE's roughly 52% reflects the opposite path, an earlier and deeper shift toward private and employer-based financing. HSTP's stated 2030 targets would move Saudi Arabia further toward the UAE's end of that range without fully adopting it, which frames the reform less as a choice between a state-run and a privatized model and more as a question of how far along an already-established regional spectrum the country moves, and how quickly.

Alsubahi, Pavlova, Alzahrani, Ahmad, and Groot's (2024) systematic review of patient-perceived healthcare quality across GCC countries adds a cautionary note to this regional picture: patients across the GCC, not only in Saudi Arabia, commonly report diagnostic and medication errors, poor provider-patient communication, missed appointments, and difficulty

accessing emergency care, suggesting that care-quality problems are a regional pattern rather than something unique to any single country's reform path.

2.3 The Saudi Journey: From Public Provision to Corporatized Care

Walston, Al-Harbi, and Al-Omar (2008) documented the strains on the Saudi system as it existed prior to Vision 2030: heavy reliance on expatriate physicians, a curative rather than preventive orientation, and rising cost pressure from an aging population and a growing chronic-disease burden. Their diagnosis (a system that had succeeded at expanding basic access but was not structured to manage long-run chronic care costs) anticipated much of the rationale HSTP would later use to justify reform.

The more recent literature shifts from diagnosis to description of the reform itself.

Chowdhury, Mok, and Leenen (2021) describe the "new model of care" introduced under HSTP, detailing the transfer of hospital operations to the Health Holding Company and the reorganization of MOH facilities into regional clusters. Alasiri and Mohammed (2022) provide a broader institutional overview of healthcare transformation since Vision 2030's launch, cataloguing the program's stated goals across access, quality, and financing. Suleiman and Ming's (2025) more recent institutional review reaches a similar descriptive conclusion but adds an implementation-focused caveat largely absent from the earlier overviews. Their account foregrounds the shift toward patient-centered, individualized care, the National Guidelines Center's role in standardizing clinical protocols nationwide, and electronic health records and telemedicine as the technical backbone of the reformed model, themes that echo the digital-platform literature discussed in Section 2.5. Where the review adds something the earlier institutional overviews do not is its treatment of implementation obstacles: insufficient resources in remote and lower-income regions, a shortage of healthcare professionals willing to work outside urban centers, staff resistance rooted in "fear of the unknown" and concern over professional autonomy, an uneven digital divide limiting telemedicine reach, and the

data-security exposure that comes with rapid digitization. None of these obstacles is new to the broader reform literature individually, but Suleiman and Ming's (2025) review is unusual among the sources surveyed here for naming them together as a connected set of implementation risks rather than treating each as an isolated footnote, a useful counterweight to the more programmatic tone of the literature it otherwise resembles. Rahman and Al-Borie (2020) and Rahman (2020), writing separately, both use document analysis to trace how Vision 2030's stated healthcare objectives translate into specific policy instruments: public-private partnerships, insurance expansion, and the corporatization of MOH facilities. What this body of work shares is a largely institutional and descriptive orientation: it documents what the reform is designed to do and how its component parts fit together, generally drawing on policy documents and structured interviews with system architects rather than on longitudinal outcome data. A more critical strand of the literature complicates this picture. Alkhamis's (2017) review of the broader privatization literature finds only weak, low-validity evidence that privatization improves access to care in general, and reports that it could not identify peer-reviewed studies examining this question specifically for Saudi Arabia at the time of the review, a gap in the evidence base rather than a finding about the Saudi case itself. Alayed and colleagues' (2024) systematic review revisits the same question seven years later with a larger pool of studies available, and still reports a mixed evidence base on access and quality outcomes rather than a settled consensus in either direction. Read together, the two reviews describe a literature that has grown without converging: more studies exist now than in 2017, but they do not agree with each other any more than the earlier, thinner literature did. This is the specific point at which the existing literature stops short of empirical resolution, and where this thesis is positioned to contribute: an assessment of what the available longitudinal data show HSTP has actually done, building on a body of work that already covers what HSTP intends.

2.4 Health Insurance, Corporatization, and Out-of-Pocket Spending

The financing side of the reform has drawn a somewhat more empirical literature than the institutional side, in part because household expenditure surveys are easier to conduct than system-wide financial audits. Al-Hanawi, Mwale, and Kamninga (2020) examined how health insurance coverage affects health-seeking behavior in Saudi Arabia, finding that insured individuals are more likely to seek care when needed, a finding relevant to the reform's stated goal of expanding effective access alongside formal insurance coverage. Alzahrani, Pavlova, Alsubahi, Ahmad, and Groot's (2025) systematic review of the cooperative health insurance system's contribution to universal health coverage finds substantial gains in nominal coverage over the past decade, alongside persistent gaps in coverage depth for specific services.

Al-Hanawi and colleagues' (2020) finding that insured individuals seek care more readily carries an implication for how the out-of-pocket trend in Chapter 4 should be read, since the two findings pull in different directions if considered together. If mandatory insurance is working as the health-seeking-behavior literature predicts, more insured individuals should be visiting providers rather than deferring care, which would tend to raise total health-system utilization and, all else equal, put upward rather than downward pressure on total spending even where insurance is absorbing a larger share of each individual bill. Chapter 4's period-average finding, a roughly 12% relative decline in the out-of-pocket share of spending, is a statement about the share insurance absorbs, not about total utilization or total spending, and the two are not the same claim: out-of-pocket burden as a percentage of current health expenditure can fall even while total visits, and total spending, rise, if the growth in insured, insurance-covered utilization outpaces the growth in the smaller uninsured or partially covered residual. This thesis's data cannot separate these two effects directly, since Chapter 4's OOP series measures a share, not a level, but the distinction matters for what a declining

OOP share does and does not imply: it is evidence consistent with insurance absorbing more of the bill, not evidence that households are spending less on health care in absolute terms, and the two are easy to conflate if the percentage figure is read on its own.

The out-of-pocket expenditure literature is the most directly relevant to this thesis's fourth hypothesis, and it points to two related but distinct forms of inequality. Al-Hanawi's (2021) national decomposition finds that out-of-pocket burden in Saudi Arabia is concentrated specifically among lower-income households (concentration index = -0.151). Separately, Almalki and colleagues' (2022) study of household spending in Riyadh by number of chronic conditions, and Al-Hanawi and Njagi's (2022) Blinder-Oaxaca decomposition comparing chronically and non-chronically ill households, both find that burden also rises with chronic-disease status. Read together, these studies show that out-of-pocket burden in Saudi Arabia is not distributed evenly along either income or health-status lines, with poorer households and chronic-disease households, the latter precisely the population Vision 2030's health reforms are meant to serve better, each bearing a disproportionate share. This is a direct empirical link to the aging-and-chronic-disease pressure identified in Section 2.1 as one of the two structural drivers of the reform in the first place.

A related but separate policy, affecting insurance coverage rather than out-of-pocket burden directly, belongs alongside this literature: the fiscal-consolidation measures introduced under the government's Fiscal Balance Program: Balanced Budget 2020 (Arab News, 2016).

Among these was a monthly levy on expatriate workers' dependents, phased in from July 2017 at SR100 per dependent and rising SR100 annually to SR400 by 2020, confirmed at the time by Finance Minister Mohammed Al-Jadaan. This is a fiscal and labor-market policy rather than a healthcare-specific one, and it does not appear in the healthcare-reform literature reviewed above, but it bears directly on Chapter 4's insurance analysis because CCHI coverage in Saudi Arabia is tied overwhelmingly to expatriate formal employment, and a

policy that raises the cost of retaining expatriate dependents in the Kingdom plausibly affects the size of the insured expatriate population independently of anything the health insurance reform itself has done. Chapter 4, Section 4.4 returns to this point in evaluating the CCHI enrollment series.

2.5 The Digital Frontier: Telehealth, Sehaty, and the Digital Divide

Sehaty and Nafis are frequently cited in Saudi policy literature as evidence of the health system's digital-first orientation, but the empirical literature on their actual reach is more cautious than the policy framing. Meraya and colleagues (2025) surveyed 466 older Saudi adults with chronic conditions in the Jazan region on telehealth use; consistent with Saudi Arabia's near-universal mobile network coverage, including in rural areas (the national regulator reports mobile voice and internet coverage reaching 99% of villages and hamlets; Communications, Space and Technology Commission, 2024), the broader Saudi telehealth-adoption literature attributes non-use less to outright lack of infrastructure than to digital literacy gaps, device access, and unfamiliarity with the platforms themselves (Alodhialah, Almutairi, & Almutairi, 2024). Alqahtani and Almaghaslah's (2025) cross-sectional study in the Aseer region similarly found mixed elderly attitudes toward mobile health applications, reinforcing that the gap is substantially a literacy and trust gap rather than a purely infrastructural one.

This matters for how the "digital divide" hypothesis in this thesis should be framed. The original assumption, namely that rural populations lack the infrastructure to use Sehaty and Nafis, is only partly supported by the literature; Saudi Arabia's mobile network coverage does not, in fact, stop at the city limits. The more precise and better-evidenced claim is that usage gaps track age and digital literacy at least as strongly as geography, and that urban and rural elderly populations may face a similar barrier for different underlying reasons. Chapter 4 evaluates this more precise version of the equity hypothesis against this literature; as Chapter

3 states, no new region-disaggregated dataset was assembled for this thesis, so this remains a literature-based reframing rather than an independently tested empirical result.

2.6 Theoretical Framework

This thesis interprets its findings through three complementary lenses, following the same triangulated structure proposed in the original research design. The four hypotheses set out in Chapter 1, efficiency, quality, equity, and financial protection, map closely onto the performance dimensions in the "control knobs" health-reform evaluation framework developed by Roberts, Hsiao, Berman, and Reich (2003), which evaluates reforms along financing, payment, organization, regulation, and behavior as policy levers against efficiency, quality, access/equity, and financial-risk-protection outcomes. This thesis works backward from HSTP's own stated claims rather than forward from that pre-specified policy-lever taxonomy, and does not attempt to classify HSTP's specific reforms against all five levers; the framework is acknowledged here because it offers an established vocabulary for exactly the outcome dimensions this thesis tests, and a reader familiar with it will recognize the structural resemblance on the outcome side even though this thesis does not adopt its lever-side analysis.

2.6.1 Theory of Change

A Theory of Change frames HSTP as a causal chain running from inputs (reallocated public revenue, royal decrees, the creation of the HHC) through activities (corporatizing hospital operations, launching digital platforms, shifting toward value-based reimbursement) to outputs (Health Clusters, an expanded insurance network) and, ultimately, to outcomes (lower government health expenditure, improved patient satisfaction, reduced out-of-pocket burden, closed urban-rural gaps). The analytical value of this framework is that it separates the question of whether the reform has produced its intended *outputs* (largely settled: the

institutional literature reviewed in Section 2.3 confirms the Health Clusters and insurance network exist and are operating) from the harder question of whether those outputs are producing the intended *outcomes*. This thesis is concerned almost entirely with the second question.

2.6.2 Agency Theory and Market Failure in Health Economics

Arrow's (1963) foundational analysis of medical care as an economic good identified the structural feature that makes healthcare markets behave differently from most others: severe information asymmetry between the physician (agent) and the patient (principal), compounded by uncertainty about both the incidence of illness and the effectiveness of treatment. Because the patient generally cannot independently verify whether a recommended test or procedure is medically necessary, a purely profit-motivated provider has a standing incentive toward induced demand, providing more or costlier care than the patient's condition requires. Value-based care, the reimbursement model HSTP is explicitly moving toward, is a direct policy response to this problem: it attempts to realign the provider's incentive with the patient's outcome rather than with service volume. Whether that realignment is actually occurring, or whether providers under the new model continue to behave as the classic principal-agent problem predicts, is an empirical question this thesis addresses indirectly through spending and satisfaction trends, given that the granular utilization data needed to test agency-theoretic predictions directly are not part of the public record (see Chapter 3).

2.6.3 Social Determinants of Health and Equity

The WHO Commission on Social Determinants of Health's (2008) final report, and the earlier "Solid Facts" synthesis by Wilkinson and Marmot (2003), establish that health outcomes are shaped substantially by conditions outside the clinical encounter: income, geography, education, and access to infrastructure. Applied to this thesis, the SDOH framework is what

motivates treating equity as a first-order outcome rather than a secondary concern to efficiency. A reform that reduces government health expenditure while widening the gap between Riyadh and the Northern Borders region would, by this framework, represent only a partial success: efficient in the narrow fiscal sense, but a failure against the "Vibrant Society" pillar's own stated commitment to a health system that serves the whole Kingdom.

Table 2.2: Theoretical Framework Application Map

Lens	Introduced	Applied in	Directness of application
Theory of Change	§2.6.1	§5.2 (output/outcome distinction)	Direct: structures the central interpretive move of the discussion chapter
Agency Theory	§2.6.2	§5.2 (induced-demand discussion)	Indirect: the data cannot test the mechanism directly; used to interpret spending and satisfaction trends as consistent or inconsistent with the theory, not to test it
Social Determinants of Health / Equity	§2.6.3	§4.3, §5.2	Direct on the equity hypothesis's framing; not used to analyze the quantitative Chapter 4 findings for H1, H2, or H4

All three lenses inform how the findings are interpreted in Chapter 5, but none of them structure the quantitative analysis in Chapter 4, which proceeds by testing each hypothesis against its own data source independently of the theoretical framework. Chapter 2's theoretical framework functions as an interpretive lens applied after the fact, not as an analytical engine that generated Chapter 4's methodology.

2.7 Synthesis and Gap Statement




Figure 2.1: Literature Base by Publication Year (n=22 academic sources)

The publication-year distribution of this thesis's academic sources (Figure 2.1) makes the timing point above concrete: 14 of the 22 academic sources cited (roughly two-thirds) were published in 2020 or later, clustered in the HSTP-implementation era rather than the earlier 2016–2019 design period. This confirms that the descriptive and institutional literature has caught up with the reform since it launched, reinforcing the case against the "under-studied" framing below, while also showing the literature base is concentrated in a narrow recent window rather than reflecting decades of accumulated scholarship on Saudi health reform specifically.

Four observations follow from this review. First, the descriptive and institutional literature on HSTP is substantial and growing: the reform's design, its component institutions, and its stated objectives are well documented, and the earlier characterization of this as an under-studied topic is, on the evidence gathered here, an overstatement of the actual gap. Second, the empirical literature that does exist is fragmented by topic: out-of-pocket spending has been studied with some rigor, digital health adoption has been studied specifically among older adults, but no existing study brings government expenditure trends, out-of-pocket

spending trends, patient satisfaction evidence, and regional equity data together against a single set of hypotheses covering the 2016–2026 period as a whole. Third, and following directly from the digital divide literature in Section 2.5, several of this thesis's original hypotheses need to be stated more precisely than the initial proposal framed them, particularly the claim that rural populations are excluded by infrastructure, which the evidence does not support as cleanly as assumed. Fourth, the regional health-economics literature reviewed in Section 2.2 (AlRuthia et al., 2025) points to an institutional-capacity explanation for cost and pipeline outcomes, specifically the absence of systematic health-technology-assessment capacity across the GCC, that sits one level upstream of the financial and behavioral explanations the rest of this literature offers; this thesis treats it as a plausible contributing factor behind the Section 4.1 PPP shortfall rather than a confirmed mechanism, since testing it directly would require access to HSTP's internal project-appraisal process rather than the public expenditure and pipeline data this thesis's design can reach.

This last observation matters for how the efficiency hypothesis should ultimately be read. A reader could take Chapter 4's finding, that government health expenditure has not fallen and the PPP pipeline is behind its own stated pace, as evidence that HSTP's underlying reform logic is flawed. The institutional-capacity literature suggests a narrower, less sweeping alternative: the reform's logic (shift delivery to autonomous clusters, bring in private capital through PPPs, let insurance absorb financing risk) may be sound in its own terms while the specific administrative machinery needed to execute it, technology assessment being one documented example, has not yet caught up with the pace the reform's own targets assume. Chapter 4 and Chapter 5 return to this distinction directly, since it changes what a null or negative finding in this thesis's data should be read as evidence of.

The chapters that follow test the four hypotheses set out in Chapter 1 against this literature and the secondary data sources detailed in Chapter 3, with that more precise framing in place.

Chapter 3: Methodology

3.1 Research Philosophy and Design

This thesis adopts a pragmatic research philosophy: the choice of method follows from what the research questions require and what data can actually answer them, rather than from allegiance to a single epistemological tradition. The original research design for this project called for an explanatory sequential mixed-methods approach, combining a large-scale citizen survey with semi-structured interviews of health-system executives. That design was set aside for this thesis because the fieldwork it requires (a stratified survey of 800 to 1,000 respondents and 15 to 20 elite interviews) falls outside the scope of what this project can execute, and because a credible substitute exists: a substantial amount of the evidence that primary fieldwork would have produced is already available in public data and published research. The substitution is closer for some parts of the original design than others. The quantitative strand, testing expenditure, out-of-pocket, and coverage trends against published statistics, maps onto longitudinal public data reasonably well, since national statistical agencies and international bodies already collect and publish the relevant series annually. The qualitative strand does not substitute as cleanly: published interviews with system architects exist in the institutional literature reviewed in Chapter 2, but they were conducted for different purposes, by different researchers, using different question protocols, and none of them were designed to surface the operational friction or informal resistance a purpose-built interview study would be positioned to find. Chapter 5, Section 5.7 returns to this asymmetry directly when scoping what a follow-up primary-data study would need to add.

The design used here is a **secondary data and document analysis**, combining two strands. The first is a quantitative trend analysis of longitudinal public statistics on government health expenditure, out-of-pocket spending, and insurance coverage. The second is a qualitative synthesis of published research and institutional documents (annual reports, government

statements, and peer-reviewed studies) used in place of original interviews to capture the "why" behind the trends the first strand identifies. This mirrors the logic of the original two-phase design (quantitative first, qualitative second, integrated through triangulation) without requiring new data collection.

3.2 Data Sources

This design's scope is fixed by what a desk-based study can access without collecting new data directly from people: no survey was fielded, no interview was conducted, and no individual-level record of any kind was obtained or handled. Every figure in Chapter 4 traces back to a publicly published statistic, a publicly available government or institutional document, or a peer-reviewed study already in the public record, and Table 3.1 lists each source category against what it actually provides and where its limits fall. This scope decision, not a change of research question, is what separates this thesis from the mixed-methods design originally proposed in Section 3.1: the quantitative and document-analysis strand of that design survives largely intact, while the fieldwork strand does not.

Table 3.1 lists the sources confirmed to be publicly accessible during the research phase of this project (see Chapter 2 for the associated literature), together with what each source actually provides and its limitations.

Table 3.1: Secondary data sources

Source	What it provides	Coverage	Access
Ministry of Health Statistical Yearbook	Demographics, health resources (beds, facilities, staffing), service utilization	Annual, published through 2024 (2023 in PDF, 2024 in Excel)	moh.gov.sa, open
HSTP Annual Report (Vision 2030 Delivery Unit)	Program-level KPIs, narrative progress against HSTP targets	Annual since program launch (2021); 2023 edition confirmed	vision2030.gov.sa, open
Council of Health Insurance (CHI) Annual Report	Insurance coverage rates, market structure, regulatory activity	Annual; 2023 edition confirmed	chi.gov.sa, open
World Bank Open Data (World Development Indicators)	Current health expenditure (% GDP), health expenditure per capita, domestic	Annual, through 2023	data.worldbank.org, open

	government health expenditure (% GDP), out-of-pocket expenditure (% of current health expenditure)		
GASTAT (General Authority for Statistics)	National Health Survey, Health Determinants Statistics, healthcare coverage rates	Periodic; 2024 Health Determinants publication confirmed	stats.gov.sa, open
Peer-reviewed literature on out-of-pocket spending, patient satisfaction, and telehealth adoption	Household-level and regional evidence not captured in national aggregates	Various, 2020–2025	Indexed journals (see Chapter 2 references)
World Bank Open Data: consumer price index and GDP deflator (used for the Chapter 4, §4.1 real-terms robustness check)	Price-level series for real-terms deflation of nominal expenditure figures	Annual, through 2023	data.worldbank.org, open; see Appendix A, Table A.1
Arab News (2016) reporting on the Fiscal Balance Program's expatriate fee schedule	Independently retrievable primary reporting on the 2017–2020 dependent-levy schedule, sourced to the government's own program document and Ministry of Finance statements	Single article, published December 2016	arabnews.com, open

One category of data assumed available in the original research design is confirmed **not** to be publicly accessible: facility- or cluster-level financial statements for the Health Holding Company and individual regional Health Clusters. Searches conducted during the research phase located organizational and structural information about the HHC (its 2022 founding decree, capital structure, and cluster organization) but no disaggregated financial disclosures. This is a genuine data limitation: Hypothesis 1 can be tested against aggregate government health expenditure trends but not against Health Cluster-level efficiency data specifically.

3.3 Document and Data Inclusion Criteria

Sources are included where they meet three criteria: they are (1) publicly accessible without a paywall or institutional subscription, (2) either produced by a Saudi government body, an international body reporting Saudi-specific data (World Bank, WHO), or a peer-reviewed

study with a Saudi Arabian sample, and (3) dated within or bridging the 2016–2026 study window, with earlier sources (e.g., Walston et al., 2008) retained only where they establish baseline conditions the reform was designed to address. Sources describing GCC-wide patterns are included where they provide comparative context, but findings are not generalized from other GCC states to Saudi Arabia without a Saudi-specific data point to anchor them.

3.4 Analytical Approach by Hypothesis

A note on the H0/H1 language used throughout this thesis. Chapter 1 states each hypothesis as a null/alternative pair, and Chapter 4 reports a verdict for each. This framing is used as an organizing device, to state each claim in a form that could in principle be falsified, not as a claim that a formal statistical hypothesis test was conducted. With 8 non-independent annual observations per series (fewer once the study period is split into pre- and post-reform windows), this thesis does not compute a test statistic, p-value, or confidence interval for any of the four hypotheses, and no such figure should be inferred from language such as "the alternative hypothesis is not supported." What Chapter 4 reports instead is whether the observed direction and magnitude of a trend is consistent or inconsistent with each hypothesis's alternative claim, alongside the actual figures so the reader can judge the strength of that consistency directly. Where two reasonable comparison methods applied to the same series produce different readings, both are reported (see Hypothesis 4 below and Chapter 4, Section 4.4).

Hypothesis 1 (Efficiency): government health expenditure as a percentage of GDP (World Bank indicator SH.XPD.GHED.GD.ZS) and in absolute terms (MOH Statistical Yearbook, HSTP Annual Report budget figures) is compared across two fixed windows: 2016–2019 (pre-reform) and 2021–2023 (post-reform, excluding 2020 as a pandemic-driven outlier common to nearly every country's health spending in that year), using the arithmetic mean of

each window. This period-average method, rather than a trend-break or structural-break test, is used consistently for both Hypothesis 1 and Hypothesis 4 below, so that the same comparison logic is not applied selectively to whichever hypothesis it happens to favor. Because a period-average comparison alone cannot establish that corporatization caused any observed change, Chapter 4 also considers alternative explanations (oil revenue fluctuation and its effect on the GDP denominator, population growth, chronic-disease-driven cost growth identified in Chapter 2, and pandemic-related spending) before attributing any observed pattern to the reform, and states explicitly where this hypothesis's own design cannot separate a reform failure from a reform succeeding against faster-growing underlying need.

Hypothesis 1's analysis includes two further checks, reported in Chapter 4, Section 4.1 ("Robustness Check") and documented here in methodological terms. The first is a real-terms adjustment: the nominal per-capita expenditure series is deflated using Saudi Arabia's consumer price index (World Bank indicator FP.CPI.TOTL, 2010=100), rebased to a 2016 basis year, to separate the portion of the nominal per-capita rise attributable to general price inflation from the portion attributable to a real increase in resources devoted to health care. The GDP deflator (NY.GDP.DEFL.ZS) was evaluated as an alternative and set aside because it is itself driven substantially by oil-price movements over this period, which would reintroduce inside the deflator the same GDP-denominator confound the percentage-of-GDP measure already carries; CPI, which prices a household consumption basket rather than a set of GDP components, does not carry that same confound. The second check narrows the post-reform window from the original three years (2021–2023) to the two years in which the Health Holding Company was actually operational (2022–2023, following its 2022 establishment per Table 1.1), on the reasoning that 2021 predates the entity whose performance the post-reform window is meant to capture. Both checks are applied only to

Hypothesis 1's per-capita and GDP-share measures, not retroactively to Hypothesis 4, because Hypothesis 4's out-of-pocket series is already reported as a percentage of current health expenditure rather than in absolute currency terms and is not subject to the same nominal-versus-real ambiguity.

Hypothesis 2 (Quality/Satisfaction): no single national patient-satisfaction time series is publicly available at the granularity this hypothesis would ideally require. This thesis instead synthesizes findings from the published satisfaction and healthcare-quality studies identified in Chapter 2, treating convergence or divergence across independent studies as the evidentiary basis rather than a single instrument's Likert-scale output. This is explicitly a weaker form of evidence than a purpose-built survey would provide, and Chapter 4 reports it with that caveat rather than presenting synthesized findings with false precision.

Hypothesis 3 (Equity): this hypothesis is not tested against a new, independent dataset assembled for this thesis. Region-disaggregated MOH facility-distribution data exist in principle (Table 3.1) but were not extracted within this project's scope, so Section 4.3 does not report a region-by-region empirical comparison. Instead, Hypothesis 3 is treated as a literature-derived reframing: the original equity hypothesis (urban advantage driven by rural infrastructure exclusion) is evaluated against the digital-divide and telehealth-adoption literature identified in Section 2.5, which points to a literacy-and-trust gap that does not map cleanly onto geography. This is a different evidentiary basis from Hypotheses 1 and 4, and Chapter 4 labels it accordingly rather than presenting it as a fourth quantitatively tested result.

Hypothesis 4 (Insurance/Out-of-pocket spending): out-of-pocket expenditure as a percentage of current health expenditure (World Bank indicator SH.XPD.OOPC.CH.ZS) is reported two ways: as a simple endpoint comparison (2016 vs. 2023) and, to apply the same method used for Hypothesis 1 consistently, as a period-average comparison across the same 2016–2019 and 2021–2023 windows. Both are reported in Chapter 4 because they do not

agree, and reporting only one would understate how sensitive this particular finding is to the choice of comparison method for a short, volatile annual series. CHI's reported insurance coverage expansion and the household-level decomposition studies identified in Section 2.4 are used to add distributional context the national aggregate cannot provide on its own.

Table 3.2: Data Availability and Evidentiary Strength by Hypothesis

Hypothesis	Primary data	Data type	Evidentiary strength
H1 (Efficiency)	World Bank WDI; HSTP Annual Report	Longitudinal national aggregate, 8 annual observations	Moderate: real trend data, but no formal significance test and the GDP-denominator/oil-price confound is named rather than isolated (see Chapter 4 caveats)
H2 (Quality)	Published satisfaction and healthcare-quality studies	Cross-sectional, multiple independent instruments, not designed for comparison	Weak: no purpose-built national instrument; treated as an open question rather than a tested result
H3 (Equity)	Digital-divide and telehealth-adoption literature (Section 2.5)	Literature synthesis; region-disaggregated MOH data identified but not extracted	Weak: reframed based on existing evidence rather than tested against new quantitative data
H4 (Insurance/OOP)	World Bank WDI (OOP series, 8 observations); CCHI enrollment (4 non-missing years, secondary-sourced)	Longitudinal national aggregate	Moderate for the OOP series; weaker for enrollment, which is compiled from secondary industry reporting rather than a primary CCHI/CHI publication (Chapter 5, Section 5.4)

The evidentiary basis is not uniform across the four hypotheses: H1 and H4 rest on longitudinal data series, while H2 and H3 rest on literature synthesis without a matching primary dataset for this thesis's specific study period and population. Chapter 4 reports findings for all four with verdict language calibrated to this difference in strength.

3.5 Validity and Reliability

Because this design does not involve original instrument construction, the standard validity and reliability procedures for primary survey research (pilot testing, Cronbach's alpha,

forward-backward translation) do not apply. In their place, this thesis applies three safeguards appropriate to secondary-data synthesis:

12. **Source triangulation:** wherever possible, a claim is supported by more than one independent source (e.g., a government-reported figure cross-checked against an international body's independently compiled statistic, or a single study's finding checked against at least one other study before being treated as a pattern rather than an isolated result).
13. **Source currency and provenance tracking:** every statistic used in Chapter 4 is dated and attributed to its original publishing body, and the analysis flags where a "most recent" figure is in fact several years old due to publication lag, a known issue with the MOH Statistical Yearbook and GASTAT releases.
14. **Explicit non-attribution of causality:** where trends coincide with policy milestones, the text distinguishes correlation from demonstrated causation throughout, consistent with the limitation stated in Chapter 1.

The third safeguard is easier to state as a principle than to apply consistently in practice, and the CCHI enrollment discussion in Chapter 4, Section 4.4 illustrates how it operates. Several regional business-press articles (Saudi Gazette, Gulf Business, Khaleej Times, and an industry aggregator) report specific expatriate departure figures attributed to a dependent levy introduced in 2017. The full archival text of these articles was inaccessible for direct verification during the research phase; the figures they report were visible only in secondary search-result summaries, not the original reporting. Consistent with the second safeguard above, they are not reported in Chapter 4 as fact. An Arab News article describing the same policy, sourced directly to the Ministry of Finance's Fiscal Balance Program document and to on-the-record statements by the Finance Minister, is retrievable in full and is cited (Arab News, 2016). The distinction that determines inclusion is source verifiability against its own original text, not source type; a government source and a press source are held to the identical standard.

3.6 Limitations of the Design

Four limitations, introduced in Chapter 1, are elaborated here in methodological terms. First, the absence of Health Cluster-level financial data means Hypothesis 1 can be tested only at the national aggregate level, which may mask efficiency gains or losses that vary substantially by region or cluster. This same data gap is what keeps the formal-versus-real-autonomy question introduced in Chapter 1, Section 1.2 open: distinguishing a genuinely autonomous, well-run cluster from a formally autonomous but centrally controlled one would require exactly the cluster-level budget and utilization data that is not publicly available, so this thesis can report the national aggregate trend but cannot say whether that trend is uniform across clusters or concentrated in a subset of them. Second, the satisfaction evidence synthesized for Hypothesis 2 was collected by different research teams, at different times, using different instruments and sample frames; treating it as a single coherent trend risks overstating the consistency of the underlying evidence, and Chapter 4 is written to avoid that overstatement. Third, publication lag in official statistics means the most recent 1–2 years of the 2016–2026 study window are necessarily represented by fewer or more provisional data points than earlier years, which limits how confidently this thesis can characterize the very latest phase of the reform. Fourth, the real-terms deflation applied in Chapter 4, Section 4.1 uses Saudi Arabia's economy-wide consumer price index (FP.CPI.TOTL) because that is the only deflator series this thesis could independently verify through the World Bank API; a healthcare-specific price index, tracking the cost of medical goods and services rather than a general household consumption basket, would in principle be a more precise deflator for health expenditure specifically, since medical cost inflation does not necessarily move in step with the broader CPI. No such Saudi-specific medical-care price index was located as publicly available during this project, so the economy-wide CPI is used as the best available

approximation, and the real-terms figures in Table 4.4 should be read with that approximation in mind rather than as a medical-cost-specific deflation.

Chapter 4: Results and Analysis

This chapter tests each of the four hypotheses set out in Chapter 1 against the secondary data sources described in Chapter 3. Figures are drawn from the World Bank's World Development Indicators (sourced from the WHO Global Health Expenditure Database, last updated December 2025), GASTAT publications, and the peer-reviewed literature identified in Chapter 2. Where a figure could not be independently verified within the scope of this thesis, that gap is stated explicitly rather than filled with an estimate.

Two sections in this chapter go beyond a single reported finding. Section 4.1 includes a robustness check that re-tests the efficiency finding under a real-terms, CPI-deflated measure and a narrower post-reform window, alongside the original nominal, wide-window comparison. Section 4.4 discusses a better-timed alternative explanation, a documented expatriate dependent levy, for part of the CCHI insurance-enrollment decline.


4.1 Hypothesis 1: Efficiency

Table 4.1: Saudi Arabia government and total health expenditure, 2016–2023

Year	Current health expenditure (% of GDP)	Domestic general government health expenditure (% of GDP)	Health expenditure per capita (current US\$)
2016	6.24	4.34	1,353.22
2017	6.03	4.33	1,400.90
2018	5.34	3.75	1,488.98
2019	5.47	4.03	1,504.55
2020	6.65	5.20	1,574.70
2021	5.93	4.57	1,654.74
2022	5.31	4.10	1,830.89
2023	5.69	4.42	1,825.25

Source: World Bank World Development Indicators (SH.XPD.CHEX.GD.ZS, SH.XPD.GHED.GD.ZS, SH.XPD.CHEX.PC.CD), retrieved via World Bank API, December 2025 update.


Figure 4.1: Saudi Arabia Health Expenditure as % of GDP, 2016–2023



Source: World Bank World Development Indicators (SH.XPD.CHEX.GD.ZS, SH.XPD.GHED.GD.ZS).

Figure 4.1: Saudi Arabia Health Expenditure as % of GDP, 2016–2023

Figure 4.2: Saudi Arabia Health Expenditure per Capita, 2016–2023



Source: World Bank World Development Indicators (SH.XPD.CHEX.PC.CD).


Figure 4.2: Saudi Arabia Health Expenditure per Capita, 2016–2023

The Health Holding Company became operational in 2022 and HSTP itself launched in 2021, giving a rough pre/post split at 2021. Using the period-average method fixed in Chapter 3 (Section 3.4), averaging government health expenditure as a share of GDP across 2016–2019 (pre-reform) gives 4.11%; averaging 2021–2023 (post-reform, excluding the 2020 pandemic

anomaly) gives 4.36%. On this measure, government health spending as a share of GDP was slightly *higher*, not lower, in the period following HSTP's launch than in the four years preceding it. Health expenditure per capita shows a clearer trend: a rise from \$1,353 in 2016 to \$1,825 in 2023, an increase of roughly 35% in nominal terms. Saudi Arabia introduced value-added tax in 2018 and tripled it from 5% to 15% in 2020, both inside this window; part of the nominal rise reflects general price inflation rather than a real increase in the volume or quality of care purchased, so the 35% figure should be read as an upper bound on real resource growth, not a precise estimate of it. The robustness check later in this section constructs a CPI-deflated real-terms series to net out that price effect directly.

Neither pattern supports Hypothesis 1's alternative claim of a measurable reduction in government health expenditure. This finding should be read carefully rather than as a simple rejection of HSTP's efficiency case, for two distinct reasons. First, percentage-of-GDP figures are sensitive to Saudi Arabia's own GDP volatility, which is itself driven substantially by oil price movements unrelated to healthcare policy; a growing denominator can lower this ratio even where health spending is flat or rising, and a shrinking one can raise it even where health spending is falling. This is not a hypothetical concern for the specific years examined here: benchmark crude oil prices fell sharply in 2020 and then rose substantially through 2021–2022 as the global economy reopened, a pattern that plausibly affects Saudi nominal GDP independently of any health-sector policy; this thesis did not obtain a verified annual oil-price series to net this effect out directly, which leaves the denominator caveat named but not resolved. Second, and more fundamentally, this hypothesis's design cannot distinguish two very different underlying stories: a reform that is failing to control costs, and a reform that is improving efficiency per unit of care while underlying demand, driven by the population growth and chronic-disease burden identified in Chapter 2, grows even faster. Total spending can rise under either story, and the data available to this thesis cannot separate

them. The per-capita rise is consistent with population growth, general inflation, and rising treatment costs for the chronic disease burden identified in Chapter 2, independent of whether the corporatized delivery model is more or less efficient per unit of care delivered, a quantity this thesis's data sources cannot isolate.



Source: World Bank World Development Indicators (SH.XPD.GHED.GD.ZS), all six countries.


Figure 4.5: Government Health Expenditure (% of GDP), Saudi Arabia vs. GCC Peers, 2016–2023

Placed against its five GCC neighbors on the same World Bank series (Figure 4.5), Saudi Arabia's government health expenditure sits mid-pack: consistently above the UAE, Bahrain, and Qatar, close to Oman, and below Kuwait in most years. The regional pattern moves broadly together, a joint rise around 2020 followed by a partial retreat, which is consistent with a shared regional driver (most plausibly the pandemic and the 2020–2022 oil price cycle) rather than country-specific reform effects. Saudi Arabia does not stand out from its neighbors either before or after HSTP's 2021 launch, which weakens the case for reading its own trend as reform-driven rather than regional.

Two supplementary indicators bear on the efficiency question more directly, both drawn from the HSTP Annual Report 2023 (Vision 2030 Delivery Unit, 2023). First, HSTP's own stated

target is to raise the private sector's share of healthcare delivery from roughly 25–28% currently to 35% by 2030; as of the most recent published figures, that share has not yet moved meaningfully toward the target. Second, of a pipeline described as more than 100 planned public-private partnership projects worth a cumulative SAR 60–80 billion (roughly USD 13.8 billion) through 2030, only 19 projects, worth a combined USD 2.9 billion, were reported as underway as of the most recently published data (Figure 4.6). The project-count and project-value shortfalls are close to each other in proportional terms, 19 of 100+ planned projects is roughly 19% of the count, and USD 2.9 billion of USD 13.8 billion is roughly 21% of the value, which suggests the projects currently underway are not disproportionately the largest or smallest ones in the pipeline; a pipeline where only small pilot projects had advanced would show a count percentage well above the value percentage, and this thesis's data show no such gap. Chapter 2, Section 2.2 discusses one plausible structural contributor to this shortfall that sits behind the deal-level bottlenecks named in Section 5.6's recommendations: AlRuthia and colleagues (2025) find that health-technology-assessment capacity, the analytical function that would let planners weigh a project's expected value before committing capital to it, is underdeveloped across the GCC generally, Saudi Arabia included. This thesis's data cannot establish that this specific gap explains the specific 19-of-100+ shortfall reported here, but it is a documented regional pattern consistent with the direction of the finding, and it is the kind of upstream, institutional-capacity explanation a purely financial read of the pipeline numbers would miss.

Figure 4.6: Public-Private Partnership Pipeline Progress Against Stated 2030 Target



Source: HSTP Annual Report 2023 (Vision 2030 Delivery Unit). Note: linear pacing is used only to express progress relative to the stated target, not as a claim about expected completion trajectory (see §4.1 caveat).

Figure 4.6: Public-Private Partnership Pipeline Progress Against Stated 2030 Target

Large infrastructure and PPP pipelines typically follow a backloaded, S-curve completion pattern rather than a straight line, so being behind a linear pace mid-program is not, by itself, evidence of underperformance; Figure 4.6 shows the current shortfall against the stated target, not a prediction of the pipeline's eventual trajectory.

Robustness Check: Real-Terms and Window-Sensitivity Analysis

The analysis above has two open issues that bear directly on how much weight the per-capita and window-average figures can carry. First, the per-capita series in Table 4.1 is nominal: it is not adjusted for the two VAT changes that fall inside the study window, the introduction of VAT in 2018 and the increase in the rate from 5% to 15% in 2020, so part of the reported rise reflects the general price level rather than a change in the real quantity or quality of care purchased. Second, the "post-reform" window used above, 2021–2023, includes two years that sit awkwardly against the reform's own timeline: 2021, a year before the Health Holding Company existed in its current form (established 2022, per Table 1.1), and 2022, still within the broader recovery from pandemic-era healthcare disruption.

To isolate the price-level effect, the Table 4.1 per-capita series was deflated using Saudi Arabia's consumer price index (World Bank indicator FP.CPI.TOTL, 2010=100), rebased to

2016. On a nominal basis, per-capita expenditure rose 34.9% between 2016 and 2023; the same series in real, 2016-basis terms rose 20.3%. On this measure, roughly 42% of the nominal increase is attributable to the general price level rather than to a real increase in resources devoted to health care.

The GDP deflator (NY.GDP.DEFL.ZS) was considered as an alternative and rejected. Across 2016–2023 it swings between roughly 69 and 102 (2010=100), a range dominated by oil-price movements rather than consumer prices, and using it would reintroduce, inside the deflator itself, the same GDP-denominator confound Section 4.1 already flags as a limitation of the percentage-of-GDP measure above. CPI avoids that problem because it tracks the price of a household consumption basket rather than a set of GDP components sensitive to hydrocarbon export values.

On the window question, restricting "post-reform" to the two years the Health Holding Company was actually operational, 2022–2023, rather than the three-year window used above, narrows the apparent post-reform average without reversing its direction relative to the pre-reform baseline.

Table 4.4: Real-terms and window-sensitivity comparison

Measure	Pre-reform (2016–2019 avg)	Post-reform, 3-yr window (2021–2023 avg)	Post-reform, HHC-operational only (2022–2023 avg)
Gov't health expenditure (% of GDP)	4.11	4.36	4.26
Per-capita expenditure, nominal (US\$)	1,436.91	1,770.29	1,828.07
Per-capita expenditure, real, 2016 basis (US\$)	1,432.47	1,614.92	1,648.87

Source: author's calculation from the Table 4.1 underlying series and World Bank CPI indicator FP.CPI.TOTL; see Appendix B for the full year-by-year series.




Figure 4.9: Nominal vs. Real (2016-Basis) Per-Capita Health Expenditure, 2016–2023

Both checks point in the same direction as the headline finding above: expenditure has not fallen since HSTP's launch, under either the nominal or the real measure, and under either the three-year or the two-year post-reform window. What changes is the size of the apparent rise, not its direction. Once general price inflation and the pre-HHC years are removed, the increase in resources devoted to health care per person is real but more modest than the nominal, wide-window comparison suggests. This does not reverse the Hypothesis 1 finding; it sharpens it, and it is the basis for the more precisely worded verdict below.

Synthesis: The available data are inconsistent with Hypothesis 1's alternative claim, under both the nominal and the real-terms measurement approaches tested and under both the wide and the narrow post-reform window, though the size of the apparent rise narrows once inflation and the pre-HHC years are removed (see the robustness check above). Government health expenditure, on both indicators examined, has not declined since HSTP's launch, and the structural mechanism HSTP's own architects point to for future efficiency gains (private-sector delivery share, PPP investment) remains substantially behind its own stated trajectory. No formal statistical significance test underlies this finding (see Chapter 3, Section 3.4), and

it is read against two different clocks: HSTP itself is roughly at the end of its own stated five-year horizon (2021–2026, Table 1.1), not at its midpoint, while Vision 2030's overall target year of 2030 is further off. Measured against HSTP's own horizon, a null finding at this stage is a more consequential result than "too early to tell"; measured against the 2030 endpoint, a different result remains possible.

4.2 Hypothesis 2: Quality and Patient Satisfaction

No single national time series of patient satisfaction under the corporatized model is publicly available, so this hypothesis is assessed through the synthesis of published studies identified in Chapter 2 rather than through an original data source. The picture from that literature is mixed rather than uniformly positive. Institutional overview studies (Alasiri & Mohammed, 2022; Chowdhury, Mok, & Leenen, 2021) report positive self-assessments of the new model of care from within the system, but these are largely descriptive rather than measuring patient experience directly. Alsubahi et al.'s (2024) GCC-wide systematic review of patient-perceived healthcare quality found meaningful variation across countries and delivery models rather than a consistent regional improvement attributable to corporatization specifically, reporting diagnostic and medication errors, poor provider-patient communication, missed appointments, and difficulty accessing emergency care as recurring patient complaints across the region (Chapter 2, Section 2.2). None of these four problem categories is addressed by the corporatized delivery model's structural changes as such; provider-patient communication and diagnostic accuracy are functions of clinical staffing and training, not of whether a hospital is administered by a ministry department or an autonomous Health Cluster. A corporatization-focused reform could leave these specific complaints largely unchanged even if it succeeds on its own structural terms, which is a limit on what this hypothesis could show even with better data than this thesis had access to: corporatization is a lever aimed at financing and organizational structure, and the region-wide quality complaints Alsubahi and

colleagues document sit largely outside that lever's reach. The telehealth adoption studies reviewed in Section 2.5 (Meraya et al., 2025; Alodhialah et al., 2024) found satisfaction with digital services to be contingent on digital literacy and platform usability rather than uniformly high, with a meaningful share of non-users citing technical and navigational difficulty as a barrier.

GASTAT's 2024 healthcare statistics (General Authority for Statistics, 2024) offer one directly relevant data point outside the reform-specific literature: the average number of healthcare visits per resident (1.9 per year, with Saudi nationals visiting more often than non-Saudis) and a dental-visit satisfaction rate of 55.6% among those who had seen a dentist in the past year. Read plainly, the dental figure means close to half of respondents who saw a dentist were not satisfied with that specific encounter, which is a notably lower figure than the reform's quality narrative would predict for a single, fairly standardized service type; dental care does not depend on the more complex coordination-across-providers problem that value-based reform is meant to solve, so a near-even split on satisfaction with it is a modest but concrete data point against, rather than for, an assumption of broadly improving service quality. These figures describe utilization and satisfaction with specific service types rather than the health system broadly, and cannot be extrapolated into a general verdict on value-based care.

Suleiman and Ming's (2025) review, discussed in Chapter 2, Section 2.3, raises a further quality-adjacent consideration that the utilization and satisfaction figures above cannot capture: the data-security exposure that comes with the same digitization the reform's quality case depends on. Electronic health records and platform-based care coordination are central to HSTP's stated model of individualized, data-informed treatment, but a system that centralizes more patient data in more digitally accessible form also carries a correspondingly larger risk surface for data breaches and unauthorized access, a tradeoff the reform's quality

narrative does not typically foreground alongside its stated benefits. This thesis has no data on Saudi Arabia's actual health-data security incident record and does not assert that the risk has materialized. A full accounting of what "quality" means under this reform needs to weigh data-security risk against care-coordination benefit, and this thesis's satisfaction-and-utilization-focused evidence base does not do so.

Suleiman and Ming's (2025) implementation-focused review, introduced in Chapter 2, Section 2.3, adds a plausible mechanism for why a system-wide quality improvement has not clearly registered in the data reviewed above, even granting that the reform's design intends one. Their review identifies staff resistance rooted in "fear of the unknown" and concern over professional autonomy, alongside a shortage of healthcare professionals willing to work outside urban centers, as implementation-level frictions distinct from the reform's stated design. A workforce that is short-staffed in some regions and whose remaining staff are still adjusting to a changed operating model is not well positioned to deliver the more individualized, patient-centered care the reform's own architects describe as its quality objective, regardless of how sound that objective is on paper. This is a literature-based explanation for the pattern in Chapter 4's data, not a finding this thesis's own evidence can confirm directly: it is a plausible mechanism, not a tested one.

Synthesis: The evidence available to this thesis is insufficient to confirm a measurable, system-wide improvement in patient satisfaction attributable to the corporatized model. This is an open question, not a rejection of Hypothesis 2's alternative hypothesis, and is the clearest instance in this thesis of a claim that would require the primary survey instrument originally proposed (see Chapter 3, Section 3.6) to resolve with confidence.

4.3 Hypothesis 3: Equity (Urban–Rural Gap)

This thesis was not able to obtain region-disaggregated healthcare expenditure or facility-distribution figures within the scope and timeframe available; the MOH Statistical Yearbook

(Ministry of Health, 2024b) publishes this level of detail, but extracting and verifying it fell outside what this project's data-gathering phase covered. This is a limitation of the analysis, not filled in with an estimate. What this thesis can report, drawn from the digital-health literature reviewed in Chapter 2, is a more specific version of the equity question than the one originally proposed.

The original hypothesis framed the equity gap primarily as an infrastructure problem: rural populations excluded from digital-first services because the underlying network does not reach them. The evidence does not support that specific framing. Saudi Arabia has near-universal mobile network coverage, including in rural areas, as part of broader Vision 2030 digital infrastructure investment. The gap the literature actually documents is one of digital literacy, device access, and platform trust, concentrated particularly among older adults. This pattern appears in both urban and rural samples, though studies specifically situated in less urbanized regions (e.g., the Aseer region cross-sectional study) found persistently mixed attitudes toward mobile health tools even where devices and connectivity were present.

Network connectivity is not, however, the only form "infrastructure" can take, and Suleiman and Ming's (2025) review, discussed above in Section 4.2 and in Chapter 2, Section 2.3, points to a narrower but still geographic form of the equity concern the original hypothesis raised. Their review reports that "insufficient resources pose significant challenges for state institutions, particularly those situated in remote and impoverished regions," and separately identifies a shortage of healthcare professionals willing to work outside urban centers. Read alongside the digital-literacy finding above, this suggests the urban-rural equity question in Saudi healthcare is not well captured by network coverage data alone: the connectivity a rural resident has access to may be adequate, while the staffing and resourcing of the facility that connectivity is meant to connect them to is not, a distinct and more traditional form of geographic inequity than the digital-first framing this thesis originally set out to test. This

thesis cannot quantify that resourcing gap with the data available to it, the same region-disaggregated MOH data gap noted at the start of this section, so it stands as a literature-supported qualification of the reframed equity finding rather than a third, independently tested claim.

The resourcing gap also connects to the formal-versus-real-autonomy question raised in Chapter 1, Section 1.2. If Health Clusters covering remote and lower-income regions have less practical latitude, or fewer resources to exercise whatever latitude they formally hold, than clusters covering Riyadh, Jeddah, or Dammam, corporatization's promised efficiency gains and its equity effects would move in opposite directions for the same underlying reason: a governance structure that looks uniform on an organizational chart but functions unevenly in practice. This thesis's data cannot test whether cluster autonomy in fact varies this way, but it is a specific, testable version of the general resourcing concern Suleiman and Ming's review raises, and Chapter 5, Section 5.7 returns to it as a direction for future research.

Synthesis: This thesis cannot confirm or reject Hypothesis 3 as originally stated, because the geographic infrastructure disparity the hypothesis assumes is not well supported by the available evidence, while the literacy-and-trust gap that the evidence does support was not part of the original hypothesis's framing. Chapter 5 discusses this as a case where the secondary-data design surfaced a mismatch between the original proposal's assumptions and what the literature actually shows, a more useful finding than a straightforward confirmation would have been.


4.4 Hypothesis 4: Insurance and Out-of-Pocket Spending

Table 4.2: Out-of-pocket health expenditure and insurance enrollment, 2016–2024

Year	OOP expenditure (% of current health expenditure)	CCHI insured individuals (millions)
2016	14.45	~12.0


2017	14.32	—
2018	15.90	—
2019	11.40	—
2020	10.08	10.5
2021	10.16	—
2022	12.78	—
2023	14.12	12.1
2024	—	13.2

Sources: World Bank World Development Indicators (SH.XPD.OOPC.CH.ZS); insurance enrollment figures compiled from industry reporting on Council of Cooperative Health Insurance (CCHI) data.



Source: World Bank World Development Indicators (SH.XPD.OOPC.CH.ZS).

Figure 4.3: Out-of-Pocket Health Expenditure, 2016–2023




Source: Council of Cooperative Health Insurance (CCHI) enrollment figures compiled from industry market reporting; see Chapter 4 caveat and References note.

Figure 4.4: CCHI Insurance Enrollment, 2016–2024

Out-of-pocket spending follows a U-shaped pattern rather than a steady decline. It fell from 14.45% in 2016 to a low of 10.08% in 2020, before rising again to 14.12% by 2023. The 2020 trough coincides with pandemic-related restrictions on elective and non-urgent care worldwide, a utilization effect rather than a policy effect, which makes it an unreliable data point for judging the insurance mandate's independent impact. Before 2020, the series moved sharply on its own: OOP spending rose from 14.32% in 2017 to a peak of 15.90% in 2018, then fell 4.5 percentage points to 11.40% in 2019, a larger single-year move than the pandemic-era dip that follows it and one this thesis cannot attribute to a specific cause with the data available. It predates HSTP's 2021 launch, so it stands as an unexplained feature of the series rather than part of the reform-era narrative that follows.

A simple endpoint comparison of the 2016 and 2023 values (14.45% to 14.12%) suggests spending is essentially back to its 2016 level. Applying the period-average method used for Hypothesis 1 changes the picture: the 2016–2019 average is 14.02%, and the 2021–2023 average (excluding 2020) is 12.35%, an approximately 12% relative decline. The two methods disagree because 2016 and 2023 happen to land close together in value despite substantial volatility in the years between them; the endpoint comparison is blind to that interior movement, while the period average smooths over exactly which two years anchor

the comparison. Neither method is uniquely correct for a series this short and this volatile. On an endpoint reading, Hypothesis 4's alternative claim is not supported; on a period-average reading, using the identical method applied to Hypothesis 1, it is. Both readings are reported because the sensitivity of the finding to method choice is itself a result, discussed further in Chapter 5, Section 5.4. A separate caution applies regardless of which method is used: Table 4.2 measures out-of-pocket spending as a share of current health expenditure, not as a level. As Chapter 2, Section 2.4 discusses, a falling share is consistent with insurance absorbing more of the total bill even if total household health spending is flat or rising, since the share depends on how fast insured and out-of-pocket spending each grow relative to one another, not on the absolute amount a household pays. Whichever of the two verdicts above is preferred, neither should be read as a claim that Saudi households are spending less on health care in absolute currency terms; this thesis's data speak to the share insurance absorbs, not to the household budget line the reform ultimately aims to shrink.



Source: World Bank World Development Indicators (SH.XPD.OOPC.CH.ZS), all six countries.

Figure 4.7: Out-of-Pocket Health Expenditure, Saudi Arabia vs. GCC Peers, 2016–2023

Regional context (Figure 4.7) puts Saudi Arabia's out-of-pocket burden in the middle of its GCC peer group: well below Bahrain, which runs 25–32% across the entire period, moderately above Oman and Qatar, which sit in the 5–8% range throughout, and broadly comparable to the UAE and Kuwait. Saudi Arabia's U-shaped path is not mirrored region-wide; Bahrain and the UAE show their own, differently timed fluctuations, which is more consistent with each country's OOP burden responding to its own financing structure and pandemic-era health-seeking patterns than with a single shared regional shock. Bahrain's position complicates Table 2.1's citizen-care classification: Chapter 2 describes Bahrain, like Saudi Arabia, as offering largely free citizen care, yet Bahrain's out-of-pocket burden sits well above the regional range rather than below it. A citizen-facing free-care policy on paper and a low measured out-of-pocket burden in practice are not the same thing, and Bahrain's case is a reminder that the H4 finding for Saudi Arabia above is a statement about the measured national aggregate, not a claim that a stated coverage policy automatically produces a correspondingly low household cost burden.

The national aggregate figures above describe the average Saudi household, and Chapter 3 committed to triangulating them against the household-level decomposition literature identified in Section 2.4, which this section now does. Al-Hanawi's (2021) national decomposition found out-of-pocket burden concentrated specifically among lower-income households (concentration index = -0.151), a distributional pattern the national average cannot show. Almalki and colleagues (2022) and Al-Hanawi and Njagi (2022) separately found that burden rises with chronic-disease status. Neither of these studies covers the 2021–2023 period this thesis's aggregate trend addresses, so they cannot confirm or refute whether the post-2020 rise in the national OOP average was distributed evenly or fell disproportionately on poorer and chronically ill households; the aggregate trend is compatible with either possibility. What the triangulation adds is a caution against reading either the

endpoint or the period-average verdict above as a claim about typical household experience: a flat or even declining national average is consistent with a widening gap between low-burden and high-burden households, and this thesis's data cannot rule that out.

CCHI enrollment tells a similarly non-linear story, and one with an equity dimension of its own that the original equity hypothesis (framed around geography, Section 4.3) did not capture. Coverage stood at roughly 12 million individuals in 2016 (about 38% of the population at the time, with expatriate workers forming the majority of enrollees), fell to 10.5 million by 2020, and has since recovered to 12.1 million (2023) and 13.2 million (2024). Because Saudi Arabia's total population grew over the same period, this recovery has roughly kept pace with population growth rather than clearly expanding coverage as a share of the population.

The 2016-to-2020 decline is usually attributed to pandemic-driven changes in expatriate employment patterns. That attribution is incomplete on its own. Beginning in July 2017, Saudi Arabia phased in a monthly levy on expatriate workers' dependents, starting at SR100 per dependent and rising by SR100 each year to SR400 by July 2020, under the government's Fiscal Balance Program: Balanced Budget 2020 (Arab News, 2016). The fee schedule was confirmed at the time by Finance Minister Mohammed Al-Jadaan and reported before the policy took effect, meaning its start date, July 2017, predates the COVID-19 pandemic by more than two years. Regional business-press coverage reports that the levy drove a substantial decline in the expatriate population beginning well before 2020, but the specific figures in that coverage could not be independently verified against the original source text. The full archival text of the relevant pages was inaccessible during the research phase, so those figures are not repeated here as fact. This follows the same standard already applied elsewhere in this thesis to unverified secondary reporting (see Chapter 5, Section 5.4 and Appendix A). The verified point is narrower but still material: the dependent levy is a

documented policy change with a start date two and a half years ahead of the pandemic, and it offers a better-timed alternative explanation than the pandemic alone for at least the pre-2020 portion of the CCHI enrollment decline. Establishing the levy's precise share of the total decline would require primary GASTAT labor-force-survey data, which falls outside this project's scope; a pandemic-only attribution should be read as incomplete rather than as the full explanation of the 2016–2020 drop. This is a distinct metric from GASTAT's separately reported figure that 100% of Saudi citizens have "basic healthcare coverage" as of 2024: that figure reflects the pre-existing MOH provision available to citizens regardless of the CCHI insurance scheme, and the two should not be conflated as evidence of the same trend.

Expatriate residents, who make up the majority of CCHI enrollees and roughly a third of Saudi Arabia's population, do not have access to the MOH's citizen-facing basic coverage fallback; their coverage is tied to formal employment, and the 2016–2020 enrollment drop shows how quickly that coverage can contract when expatriate employment patterns shift. This thesis's equity analysis, both here and in Section 4.3, is built around a citizen-facing urban-rural axis; the citizen/non-citizen axis, where the coverage volatility documented in this section falls entirely on a population without a public fallback, is a real equity dimension this thesis did not set out to measure and does not fold into the formal H3 verdict.

Synthesis: The out-of-pocket evidence for H4 is method-sensitive rather than a single clean result: an endpoint comparison does not support Hypothesis 4's alternative claim, while a period-average comparison, using the same method applied to H1, does. This thesis does not adjudicate between the two and reports the sensitivity itself as the more informative finding, discussed further in Chapter 5.


4.5 The Life Expectancy Baseline

The four hypotheses above test HSTP against outcomes the reform itself set out to change.

One further finding, outside that formal structure, closes this chapter because it bears on how

every verdict above should be read: a discrepancy in the baseline figure behind the "Vibrant Society" pillar's own headline target. Vision 2030's own documents state a 2016 baseline life expectancy of 74 years, targeting 80 years by 2030, a stated gain of 0.43 years annually (Bah, 2018). World Bank/UN Population Division data for the same period show a materially different 2016 baseline of 78.35 years, rising to 78.73 by 2023, with a dip to 77.09 in 2021 likely attributable to pandemic mortality. More recent Saudi government and press reporting cites a current figure of approximately 78.8 to 79.7 years, a value close to the World Bank series and not to a trajectory beginning at 74.

Bah's (2018) own demographic analysis provides a comparator for the 0.43-year annual target: countries with a life expectancy baseline in the 70s, the range implied by Vision 2030's own 74-year starting point, have historically achieved average annual gains below 0.31 years, short of the stated target on its own terms and before any question about which baseline is correct. That comparator is anchored to the 74-year starting point specifically; if 78.35 is the more accurate 2016 baseline, the relevant historical comparator population and achievable gain rate would differ, since countries already in the high 70s face a different demographic ceiling than countries starting from 74. The baseline discrepancy therefore does more than create two different trajectories on the same chart: it changes which comparator population the feasibility question should even be asked against.



Sources: World Bank World Development Indicators (SP.DYN.LE00.IN); Vision 2030 official target (Bah, 2018).

Figure 4.8: Vision 2030's Stated Life Expectancy Trajectory vs. World Bank/UN Data, Same Country and Years
Figure 4.8 plots both series against the same 2016–2030 timeline. The two lines start 4.35

years apart in the same year and, on the World Bank series, life expectancy was already closer to its 2030 target in 2016 than Vision 2030's own trajectory assumes it will be even by the late 2020s. The two lines reflect two different starting points for the same measurement, not the reform running ahead of or behind schedule, which makes "on track" or "behind schedule" claims about this particular target difficult to evaluate against either series in isolation.

This thesis did not attempt to adjudicate which baseline is methodologically correct; UN Population Division, WHO, and national statistical agencies frequently produce different life expectancy estimates for the same country and year, depending on the underlying life-table methodology and treatment of the large expatriate population in Saudi Arabia's demographic data. The discrepancy itself is the finding: two figures used in the same national policy conversation, differing by roughly four years for the same baseline year, is a data-quality issue independent of anything HSTP has or has not achieved, and it complicates any

straightforward claim, by this thesis or by official reporting, that the Vibrant Society target is "on track."

This discrepancy also exposes an inconsistency in how this thesis treats Vision 2030's own reporting. Section 4.1 uses HSTP's self-reported private-sector delivery share and PPP pipeline figures (Vision 2030 Delivery Unit, 2023) as evidence, without applying the same scrutiny given here to the life expectancy baseline. There is no specific reason to doubt the private-share and pipeline figures beyond the caution any self-reported program metric warrants, and neither figure was found to conflict with an independent source the way the life expectancy baseline does. The asymmetry matters: a thesis that treats one metric as suspect only after an independent source contradicts it, and treats another as reliable only because no contradicting source was located, applies skepticism unevenly rather than as a fixed standard. Chapter 5 returns to this point.

4.6 Summary of Verdicts

Table 4.3: Hypothesis Findings Summary

Hypothesis	Alternative claim tested	Finding	Primary evidence
H1 (Efficiency)	Corporatization reduces government health expenditure	Not supported under either measurement approach tested (nominal or real-terms per-capita, wide or narrow post-reform window), though the margin narrows once inflation and the pre-HHC years are removed	Table 4.1, 4.4; Figures 4.1, 4.2, 4.5, 4.9
H2 (Quality)	Corporatized, value-based care improves patient satisfaction	Inconclusive: published evidence is too fragmented across instruments and time periods to support a system-wide verdict	Literature synthesis, §4.2
H3 (Equity)	Urban populations have a measurable advantage over rural populations	Reframed rather than confirmed or rejected: available evidence points to a digital-literacy and trust gap that cuts across geography, not a purely network-infrastructure one, alongside a narrower, literature-	Literature synthesis, §4.3

		supported resourcing gap in remote regions distinct from connectivity	
H4 (Insurance/OOP)	Mandatory insurance reduces household out-of-pocket spending	Method-sensitive: an endpoint comparison (2016 vs. 2023) does not support the claim; a period-average comparison, using the same method as H1, shows an approximately 12% relative decline that does. This thesis reports both rather than choosing one (see §4.4, Chapter 5 §5.4)	Table 4.2; Figures 4.3, 4.4, 4.7

Consistent with the no-significance-test approach set out in Chapter 3, Section 3.4, this summary describes the direction and consistency of observed trends relative to each hypothesis's alternative claim, not a statistically tested result.

Read as a set, the four verdicts share a structural feature: not one of them is a clean confirmation or a clean rejection in the way the original H0/H1 framing in Chapter 1 implicitly invites a reader to expect. H1 is a rejection that gets more precise, not more certain, once tested under stricter conditions, and it is also the verdict most directly complicated by the institutional-capacity reading introduced in Chapter 2, Section 2.2 and returned to in Chapter 5, Section 5.2: a null finding on expenditure and PPP progress is compatible with a flawed reform design or a sound one outrunning its own administrative capacity, and this chapter's data cannot adjudicate between those two readings. H2 is genuinely open, constrained by the absence of a purpose-built instrument rather than by ambiguous evidence. H3 turns out to be a different, more specific question than the one originally posed, and even that reframed question now carries a further layer, since a literacy-and-trust gap and a remote-region resourcing gap can coexist rather than compete as explanations. H4 depends on which of two equally defensible methods a reader chooses. A thesis that reported four tidy yes-or-no verdicts here would be easier to summarize but would misrepresent the actual

shape of the evidence gathered in this chapter; the qualifications above are the finding, not an apology for a less clean one.

Chapter 5: Discussion and Conclusion

5.1 Summary of Findings

Of the four hypotheses tested in Chapter 4, none received unambiguous support from the available secondary data. Government health expenditure has not declined since HSTP's launch, whether measured as a share of GDP or in per-capita terms, and a subsequent robustness check confirms this under stricter, real-terms and narrower-window conditions, though with a smaller apparent rise than the initial nominal comparison suggested (H1, alternative claim not supported under either measurement approach tested). The evidence on patient satisfaction is too fragmented across studies, instruments, and time periods to confirm or deny a system-wide quality improvement, though implementation-focused literature offers a plausible partial explanation in workforce shortages and staff resistance to the changed operating model (H2, inconclusive). The equity hypothesis, as originally framed around geographic infrastructure exclusion, is not well supported by the digital-health literature, which instead points to a literacy and trust gap that cuts across geography, alongside a narrower, literature-supported resourcing gap in remote regions distinct from network connectivity (H3, reframed rather than confirmed or rejected). Out-of-pocket spending is method-sensitive: an endpoint comparison shows no sustained decline, while a period-average comparison, using the same method applied to H1, shows an approximately 12% relative decline (H4, method-dependent rather than a single clean result). A supplementary finding, not part of the original hypothesis set, is a material discrepancy between the life expectancy baseline Vision 2030 documents cite and the baseline shown in World Bank/UN data for the same year, a discrepancy that also points to an inconsistency in how uncritically Vision 2030's other self-reported figures are treated elsewhere in this thesis. A related finding concerns the 2016–2020 decline in CCHI insurance enrollment: an expatriate dependent levy phased in from July 2017, more than two years before the pandemic, offers a better-timed

alternative account than a pandemic-only explanation. One further caution applies to the H4 finding specifically: because out-of-pocket spending is measured as a share of total health expenditure rather than as an absolute level, a falling share is compatible with flat or even rising household health spending in currency terms, provided insured spending grows faster than the uninsured residual, so neither reading of the H4 evidence should be read as a claim about household budgets directly.

Read together, these findings do not indicate that HSTP has failed. They indicate that HSTP, measured against its own stated five-year horizon (2021–2026, Table 1.1) rather than Vision 2030's more distant 2030 endpoint, has not yet produced the measurable outcomes its own stated targets describe, with one hypothesis (H4) sensitive to how the comparison is made. This is a materially different and more precise claim than either "the reform is working" or "the reform is not working."

5.2 Interpretation Through the Theoretical Framework

The Theory of Change framework introduced in Chapter 2 draws a distinction between outputs and outcomes, and that distinction is where this thesis's findings are most legible. The outputs the reform promised have, by the evidence gathered here, materialized: the Health Holding Company exists and operates regional clusters, the insurance network processes over 12 million enrollees, and digital platforms are deployed nationally. The outcomes those outputs were meant to produce (lower government spending, better patient experience, narrower urban-rural gaps, lower household costs) are where this thesis finds either no measurable movement or movement in the opposite direction from what the reform's logic predicts. In Theory of Change terms, the "logical leap" from output to outcome, which Chapter 2 identified as the place a policy pipeline is most likely to break, is exactly where the evidence in Chapter 4 is weakest.

Agency theory, per Arrow's (1963) original formulation, predicts that a for-profit or quasi-corporatized delivery system will tend toward higher utilization and cost unless deliberately counteracted by a payment structure that rewards outcomes rather than volume. This thesis's data cannot test that mechanism directly: the granular utilization data needed to detect induced demand are not part of the public record, a limitation stated in Chapter 3. But the observed rise in per-capita health spending, alongside a private-delivery share that has not yet moved toward its target, is at minimum consistent with a system in which the shift to value-based reimbursement has not yet altered provider incentives at scale. This is a weaker claim than confirmation of the agency-theoretic prediction, and it is presented as such. The real-terms robustness check in Section 4.1 refines rather than overturns this reading: once general price inflation is removed, per-capita spending has still risen by a real, not merely nominal, margin, which is the pattern agency theory would predict from a system whose payment incentives have not yet shifted, and a pattern a purely inflation-driven story would not produce on its own.

Arrow's framework is usually applied to the clinician-patient relationship specifically, but the principal-agent structure it describes also holds one level up, between the reform's designers and the staff implementing it day to day, and the implementation-focused literature discussed in Sections 2.3 and 4.2 (Suleiman & Ming, 2025) is most legible through that wider reading. A regional Health Cluster's frontline staff, uncertain about their own role security and resistant to a changed operating model, face a version of the same information and incentive gap Arrow described for physicians and patients: the central reform designers cannot fully observe or control how faithfully day-to-day practice follows the new model's intent, and staff facing "fear of the unknown" have limited incentive to close that gap quickly on their own. This is a second, organizational-level agency problem layered on top of the clinical one, and it offers one plausible account of why the outputs Section 5.1 describes as largely in

place (the Health Holding Company exists, the Health Clusters operate) have not yet clearly translated into the outcomes (quality, efficiency) those outputs were built to produce.

The Social Determinants of Health framework anticipated that equity, not just efficiency, needed to be treated as a first-order outcome. The finding in Section 4.3, that the actual equity barrier in digital health is literacy and trust rather than geography, validates that framework's emphasis on non-clinical, non-infrastructure determinants of health access, even though it complicates the specific hypothesis this thesis set out to test.

The institutional-capacity reading introduced in Chapter 2, Section 2.7 belongs alongside these three lenses as a fourth interpretive caution rather than a fourth framework in its own right: where Theory of Change separates outputs from outcomes and agency theory separates incentive alignment from spending levels, the health-technology-assessment gap documented in Section 2.2 separates a reform's stated logic from its administrative capacity to execute that logic on schedule. Applied to Section 4.1's efficiency finding, this means a null result on government expenditure and PPP progress is compatible with at least two different underlying stories, a design flaw in HSTP's reform logic, or a sound design outrunning the institutional capacity built to implement it, and this thesis's data cannot distinguish between them with confidence.

5.3 Contribution and Positioning

Chapter 2 revised this thesis's original claim that Vision 2030's health reforms are critically under-studied; a substantial descriptive and institutional literature exists, including recent work directly on the reform's political economy (Hertog, 2024) and on the comparative hospital-corporatization model it most resembles (Preker & Harding, 2003). What this thesis contributes is narrower and, this author would argue, more useful at this stage of the reform: a synthesis that tests HSTP's own stated efficiency, quality, equity, and financing claims against longitudinal public data covering 2016 through 2023–2024 (the extent of the available

series; see Chapter 1, Section 1.7), rather than describing the reform's design in isolation.

Where the data support a clear verdict, this thesis states it plainly, including where that verdict runs against the reform's own narrative. Where the data do not support a clear verdict (quality, equity, and the method-sensitive out-of-pocket finding), this thesis says so rather than filling the gap with inference the evidence does not warrant.

A narrower observation about the literature itself bears on this positioning. A meaningful share of the institutional and empirical literature reviewed in Chapter 2 traces back to a small number of overlapping author groups working on Saudi and GCC healthcare reform, which is a common feature of a still-consolidating research area rather than a flaw specific to this topic, but it does mean this thesis's literature synthesis in Sections 4.2 and 4.3 draws on a narrower set of independent research perspectives than the raw citation count might suggest. Suleiman and Ming's (2025) review, incorporated in Chapter 2, Section 2.3 and drawn on again in Chapter 4, Sections 4.2 and 4.3, is one of the more recent additions to this thesis's source base to come from outside those existing clusters, and its implementation-focused findings on workforce shortages and staff resistance were independently useful for exactly that reason: they corroborated, from a source with no shared authorship with the earlier institutional-overview literature, a pattern this thesis's own data could only gesture toward indirectly.

The four-hypothesis structure this thesis tests also has a narrower positioning against the Roberts, Hsiao, Berman, and Reich (2003) control-knobs framework introduced in Chapter 2, Section 2.6. That framework separates policy levers, financing, payment, organization, regulation, and behavior, from the outcomes they are meant to produce. This thesis evaluates the outcome side only: whether efficiency, quality, equity, and financial protection moved in the directions HSTP's own stated claims predicted. It does not attempt the harder companion exercise of tracing which specific lever produced which outcome, since that would require

internal HSTP process data, which financing rule or payment reform took effect on which date and affected which population, that sits behind institutional walls this thesis's public-data design cannot reach. The four verdicts in Chapter 4 are therefore a test of direction and magnitude against the reform's stated claims, not a causal attribution to any one of the levers a control-knobs analysis would need to isolate.

5.4 Limitations

This thesis's central limitation, discussed methodologically in Chapter 3, is that it is a desk-based synthesis of secondary data rather than a study built on original survey or interview data. Seven specific consequences of that choice bear repeating here. First, as Chapter 1, Section 1.7 and Chapter 3, Sections 3.2 and 3.6 establish, Health Cluster and Health Holding Company financial data are not disclosed below the national aggregate level. One consequence carries through to this chapter: an efficient cluster and an inefficient one are indistinguishable in the data this thesis had access to, and so are a cluster with genuine operating autonomy and one still tightly controlled from the center, the formal-versus-real-autonomy question Section 1.2 and Section 5.3 raise but this thesis's data cannot resolve. Second, the patient satisfaction evidence synthesized in Section 4.2 was not designed for this purpose and does not permit the kind of before-and-after comparison a purpose-built survey would. Third, region-disaggregated expenditure and facility-distribution data, while reportedly published in the MOH Statistical Yearbook, were not extracted and verified within this project's scope, leaving Hypothesis 3 as a literature-based reframing rather than an independently tested result. Fourth, every trend reported in Chapter 4 is a correlation over time, not a demonstrated causal effect of HSTP specifically; oil price movements, population growth, and the pandemic all plausibly confound the 2016–2023 trends analyzed here. Fifth, the H4 finding on out-of-pocket spending is sensitive to which of two defensible comparison methods is used, a sensitivity this thesis reports rather than resolves, and readers should treat

H4 as an open methodological question rather than a settled result in either direction. Sixth, this thesis's equity analysis addresses a citizen-facing urban-rural axis; it does not analyze the citizen/non-citizen axis, even though Section 4.4 shows CCHI coverage volatility falling on an expatriate population with no MOH fallback coverage, a real equity dimension this thesis identifies but does not investigate. Seventh, as Section 4.5 shows, this thesis has not applied a consistent standard of skepticism to Vision 2030's own self-reported statistics. Section 4.1 treats the private-sector delivery share and PPP pipeline figures as usable evidence; Section 4.5 flags the life expectancy baseline as unreliable. The difference follows from whether an independent, contradicting source happened to be located for a given figure, not from a standard fixed in advance. Section 5.5 below returns to this point alongside the Section 4.1 robustness check. Eighth, the CCHI/CHI insurance enrollment figures reported in Chapter 4 (Table 4.2, Figure 4.4) are compiled from secondary industry reporting rather than traced to a single primary CCHI/CHI publication. CCHI/CHI's own annual reports are publicly available for every year cited, but the specific enrollee counts could not be extracted from those documents within this project's scope; this is a limitation of the data underlying Table 4.2 and Figure 4.4, not an error in how those figures are used, and is disclosed at the point of use in Section 4.4 and in Appendix A.

5.5 Robustness and Alternative Explanations

Two findings from Chapter 4 do not change the direction of this thesis's verdicts but do change how much confidence those verdicts should carry.

The real-terms and window-sensitivity analysis in Section 4.1 confirms the Hypothesis 1 finding under a tighter and more defensible test than the original nominal, wide-window comparison: government health expenditure per capita has still not fallen once general price inflation and the two years before the Health Holding Company existed are removed, but the size of the increase is meaningfully smaller, roughly 20% in real terms rather than 35% in

nominal terms. This narrows the margin by which the alternative claim is rejected without reversing the rejection itself, and it is the reason the Table 4.3 verdict for H1 is now stated as conditional on measurement approach rather than as a flat "not supported."

The expatriate dependent levy identified in Section 4.4 does something different: it does not change the H4 out-of-pocket verdict, which concerns a separate indicator, but it bears on a related pattern, the 2016–2020 CCHI enrollment decline. A pandemic-only account of that decline is incomplete; a documented fee schedule that began in July 2017, more than two years before the pandemic, offers a better-timed alternative explanation for at least part of the pre-2020 drop. This does not change the H4 verdict itself, which rests on the out-of-pocket expenditure series rather than the CCHI enrollment figures, but it changes how the enrollment figures discussed alongside H4 in Section 4.4 should be read, and it is a concrete instance of the asymmetric-skepticism problem raised in Section 4.5 and in the seventh limitation above. Read together, both findings point the same way: the headline verdicts in Chapter 4 hold up under closer scrutiny, and the closer scrutiny narrows the margin on H1 and adds a more specific mechanism to the CCHI pattern without reversing either result.

5.6 Policy Recommendations

Six recommendations follow directly from the findings, addressed to the institutions with the standing to act on them. First, the Health Holding Company and Council of Health Insurance would strengthen the evidence base for their own reform, and allow future researchers to test efficiency claims that this thesis could not, by publishing cluster-level financial and utilization data at least in aggregated, de-identified form. This would also speak directly to the formal-versus-real-autonomy question raised in Chapter 1, Section 1.2 and Chapter 2, Section 2.2 (Preker & Harding, 2003): a Health Cluster with genuine operating autonomy should show measurably different cost and utilization patterns from its neighbors over time, while clusters still tightly controlled from the center would look more uniform, and that

variation, or the lack of it, is currently invisible in the public record. Second, given that the digital divide finding in this thesis points to literacy and trust rather than infrastructure, resources currently allocated to network expansion for digital health equity would likely have greater impact if redirected toward structured digital literacy programs targeted at older adults, delivered through existing primary care touchpoints rather than through the digital platforms themselves. Third, the discrepancy identified in life expectancy baselines (Section 4.5) warrants a public reconciliation between the figure used in Vision 2030's own target-setting and the figure reported by international bodies tracking the same country and year, since the credibility of the "Vibrant Society" pillar's flagship metric depends on it being measured consistently. Fourth, the PPP pipeline shortfall identified in Section 4.1 (19 of 100+ planned projects, USD 2.9 billion of a stated USD 13.8 billion, Figure 4.6) is this thesis's most concrete and quantified efficiency-relevant finding, and it warrants a direct diagnostic rather than continued monitoring alone: specifically, a review of deal-structuring terms, investor risk perception, and regulatory or land-approval bottlenecks behind the current pace, addressed to the institutions positioned to run one. Fifth, the CCHI enrollment volatility discussed in Section 4.4 and revisited in Section 5.5, now understood to be driven partly by the expatriate dependent levy rather than by the pandemic alone, warrants monitoring specifically as an equity indicator: because expatriate residents have no MOH fallback coverage the way citizens do, a policy-driven contraction in enrollment is a coverage-loss event for that population in a way it is not for citizens, and tracking enrollment alongside expatriate population and levy-rate data separately would let regulators distinguish a policy-driven coverage dip from a genuine gap in the insurance market's functioning. Sixth, and following from the regional health-economics-capacity gap identified in Chapter 2, Section 2.2 (AlRuthia et al., 2025), the Health Holding Company and Ministry of Investment would strengthen the evidentiary basis for future PPP project selection by building or contracting a

basic health-technology-assessment function, sized to Saudi Arabia's own project pipeline rather than duplicating the systematic-review infrastructure a smaller health system would need, so that project appraisal can weigh expected value ahead of the deal-structuring stage rather than only being diagnosed after the fact, which is what recommendation four above would otherwise remain limited to doing.

Table 5.1: Policy Recommendations

#	Recommendation	Owner	Mechanism	Rough scale	Timeframe
1	Publish cluster-level financial and utilization data (aggregated, de-identified)	Health Holding Company; Council of Health Insurance	Add as a standing item in the existing annual reporting cycle	No new capital spending	Next annual reporting cycle
2	Redirect digital-equity investment from network expansion toward digital-literacy programs for older adults	Ministry of Health primary care network; Health Holding Company	Community health worker-led literacy sessions delivered at existing primary care visits	Small relative to current digital infrastructure spending	1–2 years to pilot and scale
3	Public reconciliation of the life-expectancy baseline discrepancy (Section 4.5)	General Authority for Statistics; Vision 2030 Delivery Unit	Joint technical note explaining the methodological basis for each figure	Minimal: analytical and communications effort only	Before the next HSTP annual report
4	Diagnostic review of the PPP pipeline shortfall (Section 4.1, Figure 4.6)	HSTP Delivery Unit; Ministry of Investment	Structured review of deal terms, investor risk perception, and regulatory/land-approval bottlenecks	Diagnostic study; no new capital commitment required	6–12 months
5	Track CCHI enrollment alongside expatriate population and dependent-levy rate data, reported separately from citizen coverage statistics	Council of Health Insurance; General Authority for Statistics	Add expatriate population and levy-rate columns to existing CCHI enrollment reporting	Minimal: uses data both institutions already collect	Next annual reporting cycle
6	Build a health-technology-assessment function sized to Saudi Arabia's PPP pipeline	Health Holding Company; Ministry of Investment	In-house unit or standing external contract to appraise expected	Moderate: staffing and contracting cost, no capital spending	1–2 years to establish

	(Section 2.2)		project value ahead of deal structuring		
--	---------------	--	---	--	--

These recommendations are scoped to what this thesis's own evidence supports; none requires new legislation or a change in the reform's underlying design, only additional transparency, a redirected (not expanded) budget line, and two diagnostic or monitoring additions to existing reporting.

5.7 Directions for Future Research

This thesis was explicitly designed as a substitute for the primary-data study originally proposed, not a replacement for it. Three extensions would meaningfully strengthen the findings here. A structured, nationally representative patient satisfaction survey, run at consistent intervals, would resolve the inconclusive verdict on Hypothesis 2 far more precisely than literature synthesis can. Semi-structured interviews with Health Cluster managers and CHI regulators (the qualitative strand of the original research design) would surface the operational and administrative frictions that longitudinal statistics, by construction, cannot capture, and would be the most direct way to test the formal-versus-real-autonomy question raised in Chapter 1, Section 1.2 and Chapter 2, Section 2.2. A structured comparison of budget and staffing decision rights actually exercised at the cluster level, set against the decision rights formally assigned to clusters on paper, would show directly whether the Health Holding Company's governance structure functions as Preker and Harding's (2003) framework predicts corporatized entities often do not, retaining substantive control at the center behind a devolved organizational chart, or whether Saudi Arabia's specific implementation has avoided that pattern. This thesis's public-data design cannot distinguish these two possibilities, and Section 4.1's efficiency findings are consistent with either one. Finally, a region-disaggregated analysis using the MOH Statistical Yearbook's full

geographic breakdown, which this thesis was not able to complete within its scope, would give Hypothesis 3 a proper empirical test rather than the partial, literature-based treatment offered here, and could be combined with the cluster-level autonomy comparison above to test directly whether clusters covering remote or lower-income regions exercise less practical autonomy than clusters covering major cities, the specific mechanism Section 4.3 identifies as a plausible but untested link between the equity and efficiency findings.

A fourth extension follows from the health-economics-capacity gap AlRuthia and colleagues (2025) identify across the GCC and discussed in Chapter 2, Section 2.2. If the PPP pipeline shortfall documented in Section 4.1 reflects, even partly, the absence of the health-technology-assessment capacity that would let planners weigh a given project's expected value before committing to it, then a direct study of HSTP's project-appraisal process, specifically whether and how cost-effectiveness or budget-impact analysis feeds into which PPP projects move from the planned pipeline to the underway list, would test a mechanism this thesis can only gesture toward with aggregate spending data. That kind of process-tracing study sits closer to public administration than to the health-expenditure trend analysis this thesis conducts, and would need access to internal HSTP or Ministry of Investment deliberations this thesis's public-data design cannot reach, but it would speak directly to why the pipeline is behind schedule rather than only how far behind it is.

Two narrower extensions follow directly from the robustness work in Section 5.5. A GASTAT labor-force-survey analysis of expatriate departures by year, matched against the 2017–2020 dependent-levy schedule, would let a future study estimate the levy's specific share of the 2016–2020 CCHI enrollment decline rather than leaving it, as this thesis does, as a plausible but unquantified alternative explanation to the pandemic. And extending the real-terms deflation applied in Section 4.1 back before 2016 and forward past 2023, once additional years of CPI and expenditure data are published, would show whether the

narrower, real-terms rise in per-capita expenditure documented here is a stable post-2016 pattern or specific to the particular window this thesis had data for.

5.8 Conclusion

Vision 2030's health reforms are attempting something difficult: converting a system built on unconditional entitlement into one built on regulated, insurance-mediated, partly privatized provision, without either bankrupting the state or abandoning the citizens the old system was built to serve. A decade in, the institutional machinery for that conversion (the Health Holding Company, the Health Clusters, the expanded insurance network, the digital platforms) is in place. What this thesis finds is that the outcomes that machinery was built to produce have not yet clearly appeared in the public data: government spending has not fallen even once inflation and the reform's own operational start date are accounted for, out-of-pocket costs show an ambiguous trend depending on how the comparison is made, and the equity question is not the one the reform's architects assumed it would be. None of this means the reform is failing. Measured against HSTP's own five-year implementation horizon, which is now at or near its stated end, the outcomes the reform was built to produce are not yet visible in the numbers the public can see. The robustness checks in Section 5.5, retesting the efficiency finding under stricter terms and revisiting the assumed explanation for the insurance-enrollment decline, left the direction of both findings intact. Closing the remaining gap, through the kind of primary research this thesis could not itself conduct, remains the work still ahead.

References

- Al-Hanawi, M. K. (2021). Decomposition of inequalities in out-of-pocket health expenditure burden in Saudi Arabia. *Social Science & Medicine*, 286, 114322.
- Al-Hanawi, M. K., Mwale, M. L., & Kamninga, T. M. (2020). The effects of health insurance on health-seeking behaviour: Evidence from the Kingdom of Saudi Arabia. *Risk Management and Healthcare Policy*, 13, 595–607.
- Al-Hanawi, M. K., & Njagi, P. (2022). Assessing the inequality in out-of-pocket health expenditure among the chronically and non-chronically ill in Saudi Arabia: A Blinder-Oaxaca decomposition analysis. *International Journal for Equity in Health*, 21, 192.
- Almalki, Z. S., Alahmari, A. K., Alshehri, A. M., Altowaijri, A., Alluhidan, M., Ahmed, N. J., AlAbdulsalam, A. S., Alsaiani, K. H., Alrashidi, M., Alghusn, A. G., Alqahtani, A. S., Alzarea, A. I., Alanazi, M. A., & Alqahtani, A. M. (2022). Investigating households' out-of-pocket healthcare expenditures based on number of chronic conditions in Riyadh, Saudi Arabia: A cross-sectional study using quantile regression approach. *BMJ Open*, 12(9), e066145.
- Alasiri, A. A., & Mohammed, V. (2022). Healthcare transformation in Saudi Arabia: An overview since the launch of Vision 2030. *Health Services Insights*, 15.
- Alayed, T. M., Alrumeh, A. S., Alkanhal, I. A., & Alhuthil, R. T. (2024). Impact of privatization on healthcare system: A systematic review. *Saudi Journal of Medicine & Medical Sciences*, 12(2), 125–133.
- Alkhamis, A. A. (2017). Critical analysis and review of the literature on healthcare privatization and its association with access to medical care in Saudi Arabia. *Journal of Infection and Public Health*, 10(3), 258–268.

- Alodhialah, A. M., Almutairi, A. A., & Almutairi, M. (2024). Telehealth adoption among Saudi older adults: A qualitative analysis of utilization and barriers. *Healthcare*, 12(23), 2470.
- Alqahtani, N., & Almaghaslah, D. (2025). Attitudes of elderly population toward mobile health applications in Aseer region, Saudi Arabia: A cross-sectional study. *Healthcare*, 13(19), 2464.
- AlRuthia, Y., Aldallal, S., Al-Abdulkarim, H. A., Al-jedai, A., Almudaiheem, H., Hamad, A., Elmusharaf, K., Saadi, M., Al Awar, H., Al Sabbah, H., Alghnam, S., Al Ghamdi, M., Monshi, S. S., AlAgil, N., Al Khalifa, M. E., Abdulkarim, Q., Abdulkarim, S., Jawad, H., Al-Sabahi, S., ... Alibrahim, A. (2025). Healthcare systems and health economics in GCC countries: Informing decision-makers from the perspective of the Gulf health economics association. *Frontiers in Public Health*, 13, 1510401.
- Alsubahi, N., Pavlova, M., Alzahrani, A. A., Ahmad, A., & Groot, W. (2024). Healthcare quality from the perspective of patients in Gulf Cooperation Council countries: A systematic literature review. *Healthcare*, 12(3), 315.
- Alzahrani, A. A., Pavlova, M., Alsubahi, N., Ahmad, A., & Groot, W. (2025). Impact of the cooperative health insurance system in Saudi Arabia on universal health coverage: A systematic literature review. *Healthcare*, 13(1), 60.
- Arab News. (2016, December 23). Saudi Arabia's new expat fees: What will it cost? <https://www.arabnews.com/node/1028596>
- Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5), 941–973.
- Bah, S. (2018). How feasible is the life expectancy target in the Saudi Arabian vision for 2030? *Eastern Mediterranean Health Journal*, 24(4), 401–404.
- Beblawi, H., & Luciani, G. (Eds.). (1987). *The Rentier State*. Croom Helm.

Chowdhury, S., Mok, D., & Leenen, L. (2021). Transformation of health care and the new model of care in Saudi Arabia: Kingdom's Vision 2030. *Journal of Medicine and Life*, 14(3), 347–354.

Communications, Space and Technology Commission. (2024). *Saudi internet report 2024*. Kingdom of Saudi Arabia. <https://www.cst.gov.sa>

Council of Health Insurance. (2023). *Annual report 2023*. Kingdom of Saudi Arabia. <https://www.chi.gov.sa/en/MediaCenter/ANNUAL%20REPOR/Annual%20Report%202023.pdf>

General Authority for Statistics. (2024). *Health determinants statistics publication in Saudi Arabia, 2024*. Kingdom of Saudi Arabia. <https://www.stats.gov.sa>

Gray, M. (2011). *A theory of "late rentierism" in the Arab states of the Gulf* (Occasional Paper No. 7). Center for International and Regional Studies, Georgetown University School of Foreign Service in Qatar.

Herb, M. (2014). *The wages of oil: Parliaments and economic development in Kuwait and the UAE*. Cornell University Press.

Hertog, S. (2024). The political economy of reforms under Vision 2030. In J. Sfakianakis (Ed.), *The economy of Saudi Arabia in the 21st century: Prospects and realities* (pp. 359–380). Oxford University Press.

Hood, C. (1991). A public management for all seasons? *Public Administration*, 69(1), 3–19.

Meraya, A. M., Khardali, A., Ahmad, S., Hoshabari, A. Z., Alameer, A. A., Jaafari, R. Y., Thaibah, H. A., Alfarhan, M., & Syed, M. H. (2025). Telehealth perceptions and associated factors among older adults with chronic conditions in Saudi Arabia: A comparative study of users and non-users. *Frontiers in Public Health*, 13.

Ministry of Health. (2024a). *Health Sector Transformation Program: Annual progress report*. Kingdom of Saudi Arabia.

- Ministry of Health. (2024b). *Statistical yearbook 2024*. Kingdom of Saudi Arabia.
<https://www.moh.gov.sa/en/ministry/statistics/book/pages/default.aspx>
- Preker, A. S., & Harding, A. (Eds.). (2003). *Innovations in health service delivery: The corporatization of public hospitals*. World Bank.
- Rahman, R. (2020). The privatization of health care system in Saudi Arabia. *Health Services Insights*, 13.
- Rahman, R., & Al-Borie, H. M. (2020). Strengthening the Saudi Arabian healthcare system: Role of Vision 2030. *International Journal of Healthcare Management*, 14(4), 1483–1491.
- Roberts, M. J., Hsiao, W., Berman, P., & Reich, M. R. (2003). *Getting health reform right: A guide to improving performance and equity*. Oxford University Press.
- Suleiman, A. K., & Ming, L. C. (2025). Transforming healthcare: Saudi Arabia's vision 2030 healthcare model. *Journal of Pharmaceutical Policy and Practice*, 18(1), Article 2449051.
- Vision 2030. (2016). *Kingdom of Saudi Arabia Vision 2030*. <https://www.vision2030.gov.sa>
- Vision 2030 Delivery Unit. (2023). *Health Sector Transformation Program: Annual report 2023*. Kingdom of Saudi Arabia. <https://www.vision2030.gov.sa/media/xznnx3vz/hstp-annual-report-2023-en.pdf>
- Walston, S., Al-Harbi, Y., & Al-Omar, B. (2008). The changing face of healthcare in Saudi Arabia. *Annals of Saudi Medicine*, 28(4), 243–250.
- Wilkinson, R., & Marmot, M. (Eds.). (2003). *Social determinants of health: The solid facts* (2nd ed.). World Health Organization.
- World Bank. (2025). *World Development Indicators: Health expenditure and life expectancy series for Saudi Arabia* [Data set]. <https://data.worldbank.org/indicator?locations=SA>
- World Health Organization, Commission on Social Determinants of Health. (2008). *Closing the gap in a generation: Health equity through action on the social determinants of health*. World Health Organization.

Appendix A: Data Retrieval Details

This appendix documents the exact data sources, access points, and retrieval details behind the figures reported in Chapter 4, for reproducibility. All data are drawn from publicly accessible sources; no proprietary or restricted-access datasets were used.

Table A.1: World Bank World Development Indicators

Indicator	Code	API endpoint pattern	Retrieved
Current health expenditure (% of GDP)	SH.XPD.CHEX.GD.ZS	`api.worldbank.org/v2/country/SA/indicator/SH.XPD.CHEX.GD.ZS?format=json`	December 2025 update; independently re-verified
Domestic general government health expenditure (% of GDP)	SH.XPD.GHED.GD.ZS	`api.worldbank.org/v2/country/SA/indicator/SH.XPD.GHED.GD.ZS?format=json`	December 2025 update; independently re-verified
Current health expenditure per capita (current US\$)	SH.XPD.CHEX.PC.CD	`api.worldbank.org/v2/country/SA/indicator/SH.XPD.CHEX.PC.CD?format=json`	December 2025 update; independently re-verified
Out-of-pocket expenditure (% of current health expenditure)	SH.XPD.OOPC.CH.ZS	`api.worldbank.org/v2/country/SA/indicator/SH.XPD.OOPC.CH.ZS?format=json`	December 2025 update; independently re-verified
Life expectancy at birth, total (years)	SP.DYN.LE00.IN	`api.worldbank.org/v2/country/SA/indicator/SP.DYN.LE00.IN?format=json`	Independently re-verified
Consumer price index (2010=100), used for the Chapter 4 §4.1 real-terms deflation	FP.CPI.TOTL	`api.worldbank.org/v2/country/SA/indicator/FP.CPI.TOTL?format=json`	Independently re-verified; used in the real-terms deflation reported in Table 4.4 and Appendix B

GDP deflator (2010=100)	NY.GDP.DEFL.ZS	`api.worldbank.org/v2/country/SA/indicator/NY.GDP.DEFL.ZS?format=json`	Independently re-verified; considered as an alternative deflator and not used, since it is driven substantially by oil-price movements rather than consumer prices (Chapter 3, §3.4; Chapter 4, §4.1)
-------------------------	----------------	--	---

Country code used: SA (Saudi Arabia). Underlying source for all health expenditure indicators is the WHO Global Health Expenditure Database. All values were pulled twice, independently, and matched exactly both times.

Table A.2: Institutional and government sources

Source	Access point	Content used
Ministry of Health Statistical Yearbook	moh.gov.sa/en/ministry/statistics/book	Health resource and utilization figures, Chapter 3
HSTP Annual Report 2023	vision2030.gov.sa/media/xznnx3vz/hstp-annual-report-2023-en.pdf	Private-sector delivery share, PPP pipeline figures, Chapter 4 §4.1
GASTAT Health Determinants Statistics 2024	stats.gov.sa	Healthcare coverage rate, visit frequency, dental satisfaction, Chapter 4 §4.2 and §4.4
Council of Health Insurance Annual Reports (2007–2024)	chi.gov.sa/en/MediaCenter/Pages/annual-reports.aspx	CCHI enrollment figures, Chapter 4 §4.4 and Table 4.2, compiled from secondary industry reporting rather than these primary reports directly (Chapter 5, Section 5.4).

Two categories of data referenced in the original research design were confirmed not to be publicly accessible during this project: facility- or cluster-level financial statements for the Health Holding Company and individual Health Clusters, and region-disaggregated healthcare expenditure or facility-distribution figures (which the MOH Statistical Yearbook

reportedly contains but which were not extracted within this project's scope). Both are discussed as limitations in Chapter 1 (Section 1.7), Chapter 3 (Section 3.6), and Chapter 5 (Section 5.4), and as directions for future research in Chapter 5 (Section 5.7).

Appendix B: Full Underlying Data Series

This appendix consolidates every year-by-year indicator series used in Chapter 4 into a single table. Nothing in this table is a new data point; every value already appears in Chapter 4 or was used to produce a figure or table there.

Table 4.4's real-terms comparison draws on the nominal per-capita, real per-capita, and CPI columns together; verifying that table's period-average arithmetic requires three of the eight columns below, not one. Readers working through a single hypothesis at a time may find it faster to consult the corresponding table in Chapter 4 directly (Table 4.1 for H1's core series, Table 4.2 for H4's out-of-pocket and enrollment series) and use this appendix mainly to check a value against its year, its CPI-adjusted equivalent, or a neighboring year's figure for context.

Table B.1: Consolidated indicator series, Saudi Arabia, 2016–2024

Year	Current health exp. (% of GDP)	Gov't health exp. (% of GDP)	Per-capita exp., nominal (US\$)	Per-capita exp., real, 2016 basis (US\$)	CPI, 2010=100 (FP.CPI.TOTL)	OOP exp. (% of current health exp.)	CCHI insured (millions)
2016	6.24	4.34	1,353.22	1,353.22	119.01	14.45	~12.0
2017	6.03	4.33	1,400.90	1,412.69	118.02	14.32	—
2018	5.34	3.75	1,488.98	1,465.38	120.93	15.90	—
2019	5.47	4.03	1,504.55	1,498.58	119.49	11.40	—
2020	6.65	5.20	1,574.70	1,517.28	123.51	10.08	10.5
2021	5.93	4.57	1,654.74	1,547.02	127.30	10.16	—
2022	5.31	4.10	1,830.89	1,670.37	130.45	12.78	—
2023	5.69	4.42	1,825.25	1,627.36	133.48	14.12	12.1
2024	—	—	—	—	—	—	13.2

Sources: World Bank World Development Indicators (SH.XPD.CHEX.GD.ZS, SH.XPD.GHED.GD.ZS, SH.XPD.CHEX.PC.CD, SH.XPD.OOPC.CH.ZS, FP.CPI.TOTL), retrieved via World Bank API, December 2025 update; CCHI enrollment figures compiled from industry reporting as described in Chapter 5, Section 5.4. The real per-capita column is calculated as (nominal per-capita expenditure) ÷ (CPI ÷ 119.01), rebasing the 2016 CPI value to 100; see Chapter 4, Section 4.1 (Table 4.4) for the discussion. A dash indicates no

figure was located for that year within this project's scope, consistent with the gaps already disclosed in Tables 4.1, 4.2, and 4.4.

Two gaps in this table need to be stated directly. CCHI enrollment is reported for only four of the nine years shown, reflecting the sparse and inconsistently timed nature of the industry reporting this thesis draws on for that series (Section 4.4 and Chapter 5, Section 5.4). And 2024 has no health-expenditure, CPI, or OOP figures because the underlying World Bank series had not yet been updated through 2024 as of this project's data-retrieval window; the CCHI figure for 2024 comes from a separate, more current source and does not imply the other 2024 figures exist but were omitted.