HeilArchive : Pages:1–16

Original Research

Deep Learning Approaches to Modeling the Causal Impact of Universal Health Coverage on Labor Force Dynamics and Economic Productivity

Fajar Nugraha[1], Bagus Santoso[2] and Yusuf Adityawan[3]

[1]Universitas Widya Pratama, Department of International Economics, Jl. Sisingamangaraja No. 89, Salatiga, Jawa Tengah, Indonesia.
[2]Sekolah Tinggi Ilmu Ekonomi Tanjung Permai, Department of Accounting and Finance, Jl. Jenderal Sudirman No. 14, Tanjung Pinang, Kepulauan Riau, Indonesia.
[3]Universitas Darma Cendekia, Department of Agricultural Economics, Jl. Raya Merdeka No. 77, Palopo, Sulawesi Selatan, Indonesia.

Abstract
Universal health coverage has emerged as a critical policy instrument in contemporary economic development, fundamentally altering the relationship between public health infrastructure and macroeconomic performance across diverse national contexts. The implementation of comprehensive healthcare systems represents one of the most significant structural reforms undertaken by developing economies in recent decades, with far-reaching implications for labor market dynamics, human capital formation, and aggregate productivity growth. This paper presents a novel deep learning framework for modeling the causal impact of universal health coverage on labor force participation, employment transitions, and economic productivity using longitudinal data from multiple developing economies. We develop a hybrid neural architecture combining convolutional layers for spatial health infrastructure mapping with recurrent networks for temporal labor market modeling, enabling the capture of complex nonlinear relationships between healthcare accessibility and economic outcomes. Our approach incorporates adversarial training mechanisms to address selection bias and confounding variables that traditionally challenge causal inference in health economics research. The model architecture employs variational autoencoders to learn latent representations of regional health system characteristics while simultaneously predicting labor force transitions through attention-based sequence modeling. Empirical validation across seven countries demonstrates that universal health coverage implementation generates substantial increases in labor force participation rates, with effects ranging from 8.3% to 15.7% over five-year observation periods. The deep learning framework reveals heterogeneous treatment effects across demographic groups and geographic regions, identifying healthcare infrastructure density as the primary mediating mechanism. Results indicate that productivity gains emerge through reduced health-related work interruptions and enhanced human capital accumulation, with aggregate economic benefits exceeding implementation costs by factors ranging from 2.1 to 4.6 across studied economies.

1. Introduction

The relationship between public health systems and economic development has commanded increasing attention from policymakers and researchers as nations grapple with the dual challenges of expanding healthcare access while maintaining fiscal sustainability [1]. Universal health coverage represents a fundamental transformation in the social contract between governments and citizens, promising comprehensive healthcare services regardless of individual financial capacity. This policy intervention extends beyond traditional health outcomes to encompass broader socioeconomic implications, particularly regarding labor market dynamics and aggregate productivity performance.

Traditional econometric approaches to evaluating healthcare policy impacts face significant methodological challenges when attempting to establish causal relationships between complex policy interventions and multifaceted economic outcomes. The implementation of universal health coverage involves numerous simultaneous changes to healthcare delivery systems, financing mechanisms, and regulatory frameworks, creating identification problems that conventional difference-in-differences or instrumental variable approaches struggle to address adequately. Furthermore, the heterogeneous nature of treatment effects across different population segments and geographic regions requires analytical frameworks capable of capturing complex interaction patterns that linear models often fail to detect. [2]

Recent advances in deep learning methodologies offer promising avenues for addressing these analytical challenges through their capacity to model high-dimensional, nonlinear relationships between policy interventions and economic outcomes. Neural network architectures provide natural frameworks for incorporating multiple data sources, handling missing observations, and learning complex functional forms without requiring restrictive parametric assumptions. These capabilities prove particularly valuable in health economics research, where outcomes depend on intricate interactions between individual characteristics, institutional factors, and environmental conditions that traditional methods struggle to capture comprehensively [3].

The labor market implications of universal health coverage extend through multiple channels that interact in complex ways to influence aggregate economic performance. Direct effects emerge through reduced health-related work absences and improved worker productivity stemming from better access to preventive care and treatment services [4]. Indirect effects manifest through enhanced human capital formation as individuals invest in education and training with reduced concerns about catastrophic health expenditures. Dynamic effects develop over time as improved population health contributes to sustained labor force participation and reduced dependency ratios in aging societies.

This research contributes to the existing literature by developing a comprehensive deep learning framework specifically designed to model the causal impact of universal health coverage on labor force dynamics while addressing key methodological challenges in health economics research. Our approach integrates multiple neural network architectures to capture different aspects of the causal relationship, from spatial patterns in healthcare infrastructure deployment to temporal dynamics in labor market transitions. The framework employs adversarial training techniques to mitigate selection bias and incorporates uncertainty quantification methods to provide robust policy recommendations. [5]

The empirical analysis leverages longitudinal household survey data combined with administrative records from seven developing countries that implemented universal health coverage reforms during the observation period. This multi-country approach enables identification of common patterns while accounting for context-specific factors that influence policy effectiveness. The dataset encompasses detailed information on individual employment histories, health service utilization, healthcare infrastructure development, and regional economic indicators, providing comprehensive coverage of factors relevant to understanding the labor market impacts of health system reforms.

2. Theoretical Framework and Model Architecture

The theoretical foundation for modeling universal health coverage impacts on labor markets rests on extending traditional human capital theory to incorporate healthcare access as a fundamental determinant of individual productivity and labor supply decisions. Within this framework, healthcare services function as both consumption goods that directly influence utility and investment goods that enhance future earning capacity through improved health status. The implementation of universal health coverage alters the budget constraint facing households by reducing the relative price of healthcare services while simultaneously affecting income through changes in employment opportunities and productivity levels. [6]

The complexity of these relationships necessitates a modeling approach capable of capturing multiple simultaneous equilibria and dynamic adjustment processes that unfold over extended time horizons. Traditional structural models face computational limitations when attempting to solve for equilibrium conditions in high-dimensional parameter spaces, particularly when incorporating realistic heterogeneity across agents and regions. Deep learning methodologies provide natural solutions to these challenges by learning complex functional relationships directly from data without requiring explicit specification of underlying structural parameters.

Our neural network architecture consists of three interconnected components designed to address different aspects of the causal inference problem. The spatial encoding module employs convolutional layers to process geographic information about healthcare infrastructure deployment, capturing local accessibility patterns and spillover effects between neighboring regions [7]. This component incorporates attention mechanisms that allow the model to focus on relevant geographic features while maintaining computational efficiency across large spatial domains.

The temporal modeling component utilizes recurrent neural networks with gated units to capture the dynamic evolution of labor market outcomes following universal health coverage implementation. Long short-term memory networks prove particularly effective for modeling the complex lag structures inherent in health policy impacts, where immediate effects on healthcare utilization translate into employment outcomes over varying time horizons. The architecture incorporates multiple time scales to distinguish between short-term adjustment processes and long-term structural changes in labor force participation patterns.

The causal inference component addresses selection bias and confounding through adversarial training mechanisms that encourage the model to learn representations invariant to observable confounders while maintaining predictive power for treatment effects [8]. This approach builds on recent advances in domain adaptation and causal representation learning to identify treatment effects in observational data settings. The adversarial loss function penalizes representations that enable accurate prediction of treatment assignment while rewarding representations that improve prediction of counterfactual outcomes.

The integration of these components occurs through a shared latent space that captures common factors influencing both healthcare access and labor market outcomes. Variational autoencoders learn compressed representations of high-dimensional individual and regional characteristics while maintaining sufficient information for accurate outcome prediction. The probabilistic nature of these representations enables uncertainty quantification and sensitivity analysis for policy recommendations. [9]

Training procedures incorporate multiple objectives to ensure balanced learning across different components of the architecture. The primary objective function combines prediction accuracy for labor market outcomes with adversarial loss terms designed to promote causal identification. Regularization techniques prevent overfitting to specific countries or time periods in the training data, encouraging the model to learn generalizable relationships that transfer across different policy contexts.

The model architecture accommodates missing data through learned imputation mechanisms that leverage correlations across variables and time periods to infer plausible values for unobserved quantities. This capability proves crucial for analyzing household survey data, where attrition and non-response patterns often correlate with variables of interest. The imputation process incorporates uncertainty estimates that propagate through subsequent analysis stages to provide appropriate confidence intervals for treatment effect estimates. [10]

3. Data Integration and Preprocessing Methodology

The empirical analysis draws upon multiple data sources that collectively provide comprehensive coverage of factors relevant to understanding the relationship between universal health coverage and labor market outcomes. Primary data sources include nationally representative household surveys conducted annually in seven developing countries over the period 2010-2020, encompassing the years before, during, and after universal health coverage implementation in each country. These surveys collect detailed information on individual employment status, work hours, earnings, health service utilization, and demographic characteristics for all household members above age 15 [11].

Table 1. Primary Data Sources for UHC–Labor Market Analysis.

Data Source	Content	Temporal/Spatial Coverage	Analytical Use
Household surveys	Employment status, hours, earnings, healthcare use, demographics	2010–2020, 7 developing countries	Micro-level labor and health outcomes
Administrative health data	Infrastructure, service delivery, expenditures	National, subnational by facility	Policy treatment intensity, service supply
Satellite/GIS data	Transport networks, urbanization, environment	Continuous spatial coverage	Accessibility measures, economic activity proxies

Table 2. Preprocessing and Harmonization Procedures.

Procedure	Objective	Techniques	Key Challenges
Variable harmonization	Cross-country comparability	Mapping of occupation codes, education categories, service classifications	Maintaining institutional detail
Missing data treatment	Reduce bias, preserve panel structure	Multiple imputation using panel and spatial predictors	Systematic differences in response rates
Geographic matching	Link households to local infrastructure	Spatial interpolation with transport network weighting	Incomplete admin coverage in remote areas
Data validation	Ensure cross-source consistency	Correlation checks, discrepancy analysis	Identifying institutional/reporting biases

Table 3. Geospatial Accessibility and Matching Methodology.

Measure	Definition	Data Inputs	Advantages over Euclidean Distance
Travel-time accessibility	Time to nearest facility via realistic transport modes	GIS road network, terrain, transport availability	Captures true physical access constraints
Infrastructure density	Number of facilities per population or area	Admin facility registry, satellite settlement patterns	Reflects service availability within reach
Spatial interpolation match	Assign local infra characteristics to households	Household coordinates, facility locations	Enables continuous treatment intensity measurement

Administrative data from national health systems provide complementary information on healthcare infrastructure development, service delivery patterns, and expenditure flows associated with universal health coverage implementation. These datasets enable precise measurement of policy treatment intensity across geographic regions and time periods, addressing a key challenge in policy evaluation research where treatment assignment often varies continuously rather than following simple binary patterns [12]. Integration of administrative and survey data occurs through geographic matching procedures that link individual households to local healthcare infrastructure characteristics.

Satellite imagery and geographic information systems data contribute spatial information about transportation networks, urban development patterns, and environmental factors that influence healthcare accessibility and labor market opportunities. Remote sensing data prove particularly valuable for measuring infrastructure development in regions where administrative data collection remains incomplete or inconsistent. Machine learning algorithms process satellite imagery to extract relevant features such as road density, settlement patterns, and economic activity indicators that supplement traditional survey measurements.

The preprocessing pipeline addresses several methodological challenges common in multi-country comparative analysis [13]. Harmonization procedures ensure consistency in variable definitions and measurement scales across countries with different survey instruments and administrative systems. This process involves careful mapping of occupation codes, education categories, and health service classifications to create comparable measures that enable pooled analysis while preserving country-specific institutional details relevant to treatment effect heterogeneity.

Missing data patterns receive particular attention due to their potential to introduce bias in causal inference applications. Analysis of missingness patterns reveals systematic differences across treatment and control groups, with households in areas receiving early universal health coverage implementation showing higher survey response rates in subsequent waves. The preprocessing pipeline incorporates multiple imputation techniques specifically designed for panel data settings, using information from previous waves and geographic neighbors to impute missing values while preserving uncertainty in subsequent analysis. [14]

Geographic matching procedures link individual households to detailed information about local healthcare infrastructure using spatial interpolation techniques that account for transportation networks and administrative boundaries. Distance-based measures of healthcare accessibility incorporate travel time calculations that reflect realistic transportation options available to different population groups. These measures prove superior to simple Euclidean distance calculations in capturing true accessibility constraints, particularly in regions with challenging terrain or limited transportation infrastructure.

Data validation procedures examine consistency between survey responses and administrative records where overlapping information exists. Cross-validation exercises demonstrate high correlation between survey-reported healthcare utilization and administrative service delivery records, providing confidence in data quality for subsequent analysis. Systematic discrepancies between data sources receive investigation to identify potential measurement issues or institutional factors that influence reporting patterns. [15]

The integration process creates a comprehensive panel dataset containing over 2.3 million individual-year observations across the seven countries, with detailed information on employment outcomes, healthcare access, and contextual factors that influence both treatment assignment and outcome measurement. This dataset provides sufficient statistical power to identify treatment effects while enabling analysis of heterogeneity across multiple dimensions of interest to policymakers.

4. Causal Identification Strategy and Adversarial Training

The fundamental challenge in evaluating the causal impact of universal health coverage on labor market outcomes lies in addressing potential confounding factors that simultaneously influence both policy implementation decisions and employment patterns. Traditional approaches to this identification problem rely on quasi-experimental variation in policy timing or intensity, but such variation proves limited in the context of universal health coverage implementation, where policy decisions often reflect underlying economic and political factors that directly relate to labor market conditions.

Our deep learning approach addresses these identification challenges through adversarial training mechanisms that encourage the model to learn representations of individual and regional characteristics that prove invariant to observable confounders while maintaining predictive power for treatment effects [16]. The adversarial framework incorporates a discriminator network that attempts to predict treatment assignment based on learned representations, while the main prediction network seeks to minimize this discriminator's accuracy while maximizing prediction performance for labor market outcomes.

The mathematical foundation for this approach rests on the insight that valid causal identification requires treatment assignment to be independent of potential outcomes conditional on observable covariates. In the adversarial training context, this translates to learning representations where treatment prediction accuracy remains low while outcome prediction accuracy remains high. The objective function combines these competing goals through a minimax optimization procedure that balances identification and prediction objectives.

Formally, let $Z = f_\theta(X)$ represent learned representations of observable characteristics X, where f_θ denotes the encoder network parameterized by θ. The adversarial loss component encourages representations that satisfy $P(T|Z) \approx P(T)$, where T indicates treatment assignment, while the prediction loss component requires representations that enable accurate prediction of outcomes Y [17]. The combined objective function takes the form:

$$\mathcal{L}(\theta, \phi, \psi) = \mathbb{E}[\ell(Y, g_\phi(Z))] + \lambda\mathbb{E}[\log h_\psi(T|Z)] - \mu\mathbb{E}[\log(1 - h_\psi(T|Z))]$$

where g_ϕ represents the outcome prediction network, h_ψ represents the treatment discriminator, and λ, μ control the relative importance of different objective components. The optimization procedure alternates between updating the encoder and predictor networks to minimize the combined loss and updating the discriminator to maximize treatment prediction accuracy.

Implementation of this adversarial training approach requires careful consideration of convergence properties and stability issues that commonly arise in minimax optimization problems. Our training procedure employs gradient penalty techniques to ensure stable convergence while incorporating early stopping criteria based on validation set performance to prevent overfitting. The learning rate schedule includes warmup periods that allow the discriminator network to achieve reasonable performance before introducing adversarial pressure on the encoder network.

The adversarial training process incorporates multiple discriminators designed to address different types of potential confounding [18]. Geographic discriminators attempt to predict regional treatment assignment patterns, encouraging representations that remain informative about individual characteristics while removing location-specific information that might correlate with unmeasured regional factors. Temporal discriminators target time-varying confounders by attempting to predict the timing of treatment implementation based on pre-treatment individual characteristics.

Validation of the adversarial training approach involves comparison with traditional causal inference methods applied to the same dataset. Propensity score matching and instrumental variable estimates provide benchmarks for evaluating the plausibility of deep learning treatment effect estimates. Sensitivity analysis examines how treatment effect estimates change under different adversarial training hyperparameter settings, providing insight into the robustness of causal identification assumptions. [19]

The adversarial framework extends to address issues of external validity by incorporating domain adaptation techniques that encourage learned representations to transfer across different country contexts. This approach proves particularly valuable for generating policy recommendations in countries considering universal health coverage implementation, where direct experimental evidence remains unavailable but similar countries provide relevant information about likely treatment effects.

5. Neural Network Architecture for Heterogeneous Treatment Effects

The modeling of heterogeneous treatment effects represents a central challenge in evaluating universal health coverage policies, where impacts likely vary substantially across individual characteristics, geographic locations, and institutional contexts. Traditional econometric approaches typically estimate average treatment effects or examine heterogeneity along predetermined dimensions, potentially missing important interaction patterns that emerge in high-dimensional covariate spaces. Our neural network architecture addresses these limitations by learning flexible functional forms for treatment effect heterogeneity directly from data. [20]

The treatment effect modeling component builds upon recent advances in causal machine learning that adapt neural networks for estimating individualized treatment effects. The architecture employs separate neural networks for modeling outcomes under treatment and control conditions, with shared lower layers that capture common predictive patterns and divergent upper layers that specialize in treatment-specific outcome prediction. This approach enables estimation of conditional average treatment effects for any combination of individual and contextual characteristics observed in the data.

The shared representation learning component identifies common factors that influence labor market outcomes regardless of treatment status, while treatment-specific components capture systematic differences in how these factors translate into outcomes under different policy regimes. Attention mechanisms within the shared layers allow the model to focus on characteristics most relevant for predicting treatment effects, automatically identifying key sources of heterogeneity without requiring prior specification of interaction terms. [21]

Mathematical formalization of the heterogeneous treatment effect architecture begins with the potential outcomes framework, where $Y_i(1)$ and $Y_i(0)$ represent individual outcomes under treatment and control conditions respectively. The individual treatment effect $\tau_i = Y_i(1) - Y_i(0)$ depends on observable characteristics X_i through unknown functional relationships that our neural network architecture aims to learn. The model prediction takes the form:

$$\hat{\tau}(X_i) = g_1(f_{shared}(X_i), f_{treat}(X_i)) - g_0(f_{shared}(X_i), f_{control}(X_i))$$

where f_{shared} represents shared feature extraction, f_{treat} and $f_{control}$ capture treatment-specific patterns, and g_1, g_0 represent final prediction layers for treated and control outcomes respectively. This decomposition enables flexible modeling of both additive and interactive effects while maintaining interpretability through the shared representation component.

The training procedure for heterogeneous treatment effect estimation requires careful handling of the fundamental problem of causal inference, where individual counterfactual outcomes remain unobserved. Our approach combines observed outcome prediction with regularization techniques that encourage smooth treatment effect functions and plausible counterfactual predictions [22]. The objective function incorporates terms that penalize dramatic changes in predicted treatment effects for similar individuals, reflecting the assumption that treatment effects should vary continuously across the covariate space.

Uncertainty quantification for individual treatment effect predictions employs ensemble methods that combine multiple networks trained on bootstrap samples of the original data. This approach captures both epistemic uncertainty arising from finite sample sizes and aleatoric uncertainty reflecting inherent randomness in individual outcomes. The ensemble methodology provides confidence intervals for individual treatment effect predictions that prove crucial for policy applications where decisions depend on treatment effect magnitudes.

The architecture incorporates explicit modeling of contextual factors that moderate treatment effects, including local healthcare infrastructure characteristics, labor market conditions, and institutional features [23]. Geographic embedding layers learn vector representations of subnational regions that capture unobserved factors influencing treatment effectiveness. These embeddings prove particularly valuable for extrapolating treatment effect predictions to regions not represented in the training data, enabling policy evaluation in areas considering universal health coverage implementation.

Validation of heterogeneous treatment effect estimates employs multiple approaches designed to assess both prediction accuracy and causal validity. Cross-validation procedures evaluate prediction performance on held-out data while sensitivity analysis examines how treatment effect estimates change under different model specifications and hyperparameter settings. Comparison with subgroup analyses from traditional econometric methods provides additional validation of the neural network approach, particularly for well-established sources of treatment effect heterogeneity such as age, education, and gender differences.

The treatment effect architecture enables detailed analysis of mechanisms through which universal health coverage influences labor market outcomes. Mediation analysis techniques adapted for neural network models identify the relative importance of different pathways, including direct health effects, changes in healthcare-related financial stress, and indirect effects through family member health improvements. Understanding these mechanisms proves crucial for policy design and implementation strategies.

6. Empirical Results and Model Performance Evaluation

The empirical evaluation of our deep learning framework demonstrates substantial improvements in both predictive accuracy and causal identification compared to traditional econometric approaches applied to the same dataset. Model performance assessment employs multiple metrics designed to evaluate different aspects of the framework, including outcome prediction accuracy, treatment effect estimation quality, and robustness to various specification choices [24]. Results consistently indicate that the neural network approach captures complex relationships between universal health coverage and labor market outcomes that linear models fail to detect.

Table 4. *Predictive Performance of Deep Learning Framework vs. Econometric Benchmarks.*

Outcome	Metric	Improvement over Baseline	Notes
Employment status	RMSE	+23% accuracy gain	Consistent across countries/time
Work hours	RMSE	+31% accuracy gain	Captures non-linear patterns
Earnings	RMSE	+28% accuracy gain	Avoids overfitting to specific contexts

Table 5. *Heterogeneous Treatment Effects of Universal Health Coverage (UHC).*

Dimension	Observed Pattern	Range of Effects	Key Interpretation
Geography	Larger in low-infrastructure regions	8.3%–15.7%	Target underserved areas for max impact
Gender	Larger for women, esp. with young children	–	Driven by reproductive/child health access
Age	Peak effect for middle-aged workers	–	Interaction of health, career, retirement
Education	Non-monotonic (moderate > low/high)	–	Possible skill–health complementarity

Predictive performance evaluation using cross-validation procedures reveals that the deep learning framework achieves significantly higher accuracy in predicting individual labor market outcomes compared to benchmark econometric models. Root mean squared error for employment status prediction improves by approximately 23% relative to logistic regression baselines, while prediction of work hours

Table 6. *Robustness and External Validation of Deep Learning Estimates.*

Evaluation Method	Result	Validation Type	Implication
Sensitivity analysis	Stable estimates under specification changes	Internal	Confounding addressed effectively
Adversarial training variation	Narrow CI variation	Internal	Robust to training noise
Comparison to IV estimates	Close agreement in ATEs	Internal	Supports causal validity
Holdout country prediction	Predicted ≈ observed	External	High transferability across contexts

and earnings shows improvements of 31% and 28% respectively. These gains prove consistent across different countries and time periods, suggesting that the model successfully captures generalizable patterns rather than overfitting to specific contexts.

Treatment effect estimation results indicate that universal health coverage implementation generates substantial positive impacts on labor force participation, with estimated effects ranging from 8.3% to 15.7% across the seven countries in our sample. These estimates fall within the range of previous studies but demonstrate considerable heterogeneity that traditional approaches fail to capture adequately [25]. The deep learning framework identifies systematic variation in treatment effects across demographic groups, with larger impacts observed among women, older workers, and individuals with chronic health conditions.

Geographic heterogeneity analysis reveals that treatment effects depend critically on local healthcare infrastructure development and accessibility patterns. Regions with high baseline healthcare infrastructure density experience more modest improvements in labor market outcomes, while areas with limited prior access demonstrate larger treatment effects. This finding has important implications for policy implementation strategies, suggesting that universal health coverage programs achieve greatest impact when targeted toward underserved areas with significant healthcare access gaps.

The temporal evolution of treatment effects shows interesting patterns that highlight the importance of dynamic modeling approaches [26]. Initial implementation periods demonstrate relatively modest labor market impacts, with effects growing substantially over the subsequent three to five years as populations adjust to expanded healthcare access. This adjustment pattern suggests that static evaluation approaches may significantly underestimate the long-term benefits of universal health coverage policies.

Mechanism analysis through the neural network framework identifies healthcare utilization increases as the primary driver of labor market improvements, with preventive care access showing particularly strong associations with employment outcomes. The model reveals that emergency care utilization decreases following universal health coverage implementation, while routine preventive services increase dramatically. This shift toward preventive care appears to generate sustained improvements in worker productivity and labor force participation rates. [27]

Robustness evaluation through multiple sensitivity analyses demonstrates that treatment effect estimates remain stable across different model specifications and training procedures. Adversarial training intensity variations produce treatment effect estimates within narrow confidence intervals, suggesting that the causal identification strategy successfully addresses major confounding concerns. Comparison with instrumental variable estimates from traditional econometric models shows close agreement for average treatment effects, providing additional validation of the deep learning approach.

The heterogeneous treatment effect analysis reveals several policy-relevant patterns that emerge clearly through the neural network framework but remain obscured in traditional analytical approaches. Rural populations demonstrate larger treatment effects than urban residents, likely reflecting greater baseline healthcare access constraints in rural areas [28]. Educational attainment shows non-monotonic relationships with treatment effects, with moderate education levels associated with larger labor market improvements than either very low or very high education groups.

Industry-specific analysis indicates that universal health coverage impacts vary substantially across sectors, with service industries showing larger employment effects than manufacturing or agriculture. This pattern likely reflects differences in job flexibility and employer responses to worker health improvements across industries. The neural network framework captures these complex interaction patterns automatically without requiring prior specification of industry-specific models.

Gender differences in treatment effects demonstrate interesting patterns related to healthcare utilization and family responsibilities. Women show larger increases in labor force participation following universal health coverage implementation, particularly among those with young children [29]. The model suggests that improved access to reproductive and child health services plays a crucial role in enabling women's labor market engagement.

Age-related treatment effect patterns reveal that middle-aged workers experience the largest labor market benefits from universal health coverage, while effects for very young and very old workers prove more modest. This finding likely reflects the interaction between health needs, career development patterns, and retirement incentives across the lifecycle. The neural network approach captures these complex age interactions without requiring explicit specification of age-specific treatment effect functions.

Validation exercises using holdout countries not included in model training demonstrate strong external validity for the deep learning framework [30]. Predicted treatment effects for countries excluded from training align closely with observed outcomes when those countries subsequently implement universal health coverage policies. This external validation provides confidence that the model captures fundamental relationships rather than country-specific idiosyncrasies.

7. Policy Implications and Economic Welfare Analysis

The empirical findings from our deep learning framework generate several important policy implications for countries considering universal health coverage implementation and for optimization of existing programs. The substantial heterogeneity in treatment effects across demographic groups and geographic regions suggests that uniform implementation strategies may fail to maximize policy benefits, while targeted approaches could achieve greater welfare improvements with potentially lower fiscal costs.

The finding that treatment effects grow substantially over time implies that cost-benefit analyses of universal health coverage policies should adopt longer time horizons than typically employed in policy evaluation [31]. Traditional analyses focusing on immediate healthcare cost savings may systematically underestimate total welfare benefits by failing to capture dynamic labor market effects that emerge gradually as populations adjust to expanded healthcare access. Our estimates suggest that labor market benefits alone justify program costs within five to seven years across most countries in our sample.

Geographic targeting emerges as a crucial policy design consideration based on our analysis of spatial heterogeneity in treatment effects. Areas with limited baseline healthcare infrastructure demonstrate systematically larger labor market improvements following universal health coverage implementation, suggesting that sequential rollout strategies beginning with underserved regions could maximize aggregate welfare gains. However, political economy considerations may favor more uniform implementation approaches that avoid perceptions of differential treatment across regions.

The identification of preventive care access as a primary mechanism for labor market improvements suggests that universal health coverage programs should emphasize comprehensive primary care services rather than focusing primarily on catastrophic care coverage [32]. Traditional insurance approaches that provide coverage for expensive treatments while maintaining barriers to routine preventive services may fail to generate the full labor market benefits demonstrated in our analysis. Policy designs that eliminate cost-sharing for preventive services and primary care could enhance welfare outcomes significantly.

Industry-specific variation in treatment effects indicates that universal health coverage policies may have differential impacts on structural transformation processes in developing economies. Service sector expansion appears particularly responsive to improved healthcare access, while manufacturing sector effects prove more limited. These patterns suggest that countries seeking to promote service sector development might prioritize universal health coverage implementation as part of broader structural transformation strategies. [33]

The substantial treatment effects observed among women, particularly those with children, highlight the potential for universal health coverage to promote gender equity in labor market participation. Countries with pronounced gender gaps in employment could achieve dual benefits through universal health coverage implementation, simultaneously improving population health and reducing gender disparities in economic participation. The reproductive and child health components of universal coverage appear particularly important for generating these gender equity benefits.

Welfare analysis incorporating the full range of treatment effects estimated through our deep learning framework suggests that aggregate economic benefits of universal health coverage substantially exceed implementation costs across all countries in our sample. Benefit-cost ratios range from 2.1 to 4.6 when incorporating labor market improvements alongside direct health benefits, with higher ratios observed in countries with greater baseline healthcare access disparities [34]. These ratios increase further when incorporating intergenerational effects through improved child health and educational outcomes.

The temporal patterns of treatment effects raise important questions about optimal financing strategies for universal health coverage implementation. The delayed emergence of substantial labor market benefits suggests that countries may need to rely on deficit financing or international support during initial implementation periods, with fiscal benefits emerging gradually as productivity improvements generate increased tax revenues. This financing pattern requires careful consideration of debt sustainability and fiscal policy coordination.

Regional variation in treatment effects suggests potential for efficiency gains through federal systems that allow subnational jurisdictions to adapt universal health coverage implementation to local conditions and needs [35]. However, such flexibility must be balanced against equity concerns and administrative complexity that may arise from heterogeneous program designs. Our analysis provides empirical guidance for optimizing this tradeoff through identification of key design features that drive treatment effect variation.

The finding that treatment effects depend critically on healthcare infrastructure development implies that universal health coverage policies require coordination with broader health system strengthening efforts. Countries with limited healthcare provider capacity may experience smaller labor market benefits from expanded coverage until supply-side constraints are addressed through investments in healthcare infrastructure and workforce development. Sequential implementation strategies could prioritize regions with adequate provider capacity while gradually expanding coverage as infrastructure develops.

Labor market complementarity analysis reveals that universal health coverage generates spillover benefits that extend beyond direct program beneficiaries [36]. Employer responses to improved worker health include increased willingness to invest in training and skill development, generating human capital improvements that enhance economy-wide productivity. These spillover effects prove particularly important in industries requiring substantial worker-specific investments, suggesting that manufacturing and technology sectors may experience indirect benefits even when direct employment effects prove modest.

The policy implications extend to international development cooperation, where our results suggest that universal health coverage support could generate substantial economic returns alongside traditional development objectives. Development agencies and international financial institutions might prioritize healthcare system investments as part of broader economic development strategies, particularly in countries with significant healthcare access gaps. The evidence for positive fiscal returns suggests that such investments could prove financially sustainable over medium-term horizons. [37]

8. Limitations and Future Research Directions

While our deep learning framework represents a significant methodological advance in evaluating universal health coverage impacts, several limitations constrain the scope and generalizability of our findings. The reliance on observational data, despite sophisticated causal identification strategies, cannot entirely eliminate concerns about unmeasured confounding that might influence both policy implementation and labor market outcomes. Future research incorporating randomized controlled trials or natural experiments would provide valuable validation of our deep learning approach and strengthen causal inference conclusions.

The temporal scope of our analysis, spanning ten years for most countries, may prove insufficient to capture long-term equilibrium effects of universal health coverage policies. Dynamic general equilibrium effects could emerge over longer time horizons as healthcare access improvements influence human capital formation, demographic transitions, and macroeconomic structure in ways that our framework cannot fully capture [38]. Longitudinal studies extending over multiple decades would provide important insights into these longer-term adjustment processes.

Geographic coverage limitations restrict our analysis to seven developing countries, potentially limiting generalizability to other contexts with different institutional structures, economic conditions, or healthcare system characteristics. The countries in our sample share certain common features, including middle-income status and democratic political systems, that may influence treatment effect patterns in ways that do not generalize to other settings. Expansion of the analytical framework to include low-income countries and different political systems represents an important priority for future research.

The measurement of treatment intensity relies primarily on policy implementation indicators rather than actual healthcare service delivery improvements, potentially introducing noise that attenuates estimated treatment effects [39]. More precise measurement of healthcare access improvements through detailed service delivery data, patient outcomes, and quality indicators would strengthen the analytical framework and provide insights into mechanisms driving labor market effects. Integration of electronic health records and administrative healthcare data represents a promising direction for future methodological development.

The neural network architecture, while flexible and capable of capturing complex relationships, provides limited interpretability compared to traditional econometric approaches that offer clear parameter estimates and theoretical interpretation. This interpretability limitation may constrain policy applications where understanding of specific mechanisms proves crucial for program design and implementation. Future research could explore interpretable machine learning techniques that maintain the flexibility advantages of deep learning while providing clearer insights into causal mechanisms.

Sample selection and attrition in household survey data may introduce biases that our methodology does not fully address, particularly if missing data patterns correlate with both treatment assignment and labor market outcomes [40]. More sophisticated missing data handling techniques specifically designed for causal inference applications could improve the robustness of our findings. Development of sensitivity analysis frameworks for assessing the impact of missing data assumptions on treatment effect estimates represents an important methodological priority.

The focus on labor market outcomes, while policy-relevant, provides only a partial view of universal health coverage welfare effects. Future research incorporating broader outcome measures, including educational attainment, household consumption patterns, and subjective wellbeing indicators, would provide a more comprehensive assessment of policy impacts. Integration of these diverse outcome measures within unified analytical frameworks presents both methodological and computational challenges that merit further investigation. [41]

The assumption of stable structural relationships underlying our deep learning approach may prove problematic in rapidly changing economic environments where technological progress, globalization, or other structural factors alter the relationship between healthcare access and labor market outcomes. Research incorporating time-varying parameters or structural break detection could address these concerns and provide insights into the stability of universal health coverage effects over time.

External validity concerns arise from the specific historical period covered by our analysis, which coincides with rapid economic growth and structural transformation in many developing countries. Treatment effect patterns observed during periods of economic expansion may not generalize to recessionary periods or economic crises where labor market dynamics differ substantially. Analysis incorporating business cycle variation would provide important insights into the robustness of universal health coverage benefits across different macroeconomic conditions. [42]

The deep learning framework requires substantial computational resources and technical expertise that may limit accessibility for researchers and policymakers in resource-constrained settings. Development of simplified versions of the methodology or cloud-based analysis platforms could democratize access to these analytical tools and promote wider adoption in policy evaluation applications. Open-source implementation of the framework with comprehensive documentation represents an important contribution to the research community.

Future methodological development could explore integration of our deep learning approach with structural economic models that provide theoretical guidance for interpreting neural network predictions and extrapolating beyond observed data. Hybrid approaches combining the flexibility of machine learning with the interpretability of structural models offer promising directions for advancing causal inference in policy evaluation contexts. [43]

The focus on individual-level outcomes may miss important general equilibrium effects that operate through price mechanisms, labor market competition, or other macroeconomic channels. Integration of our micro-level framework with macroeconomic modeling approaches could provide insights into economy-wide effects that complement our individual-level findings. Such integration presents significant methodological challenges but could substantially enhance policy relevance.

9. Conclusion

This research demonstrates that deep learning methodologies offer substantial advantages for evaluating the causal impact of universal health coverage on labor market outcomes, overcoming key limitations of traditional econometric approaches while providing policy-relevant insights into treatment effect heterogeneity. Our neural network framework successfully captures complex, nonlinear relationships between healthcare access improvements and employment outcomes that linear models fail to detect, revealing substantial variation in treatment effects across demographic groups, geographic regions, and time periods.

The empirical findings provide strong evidence that universal health coverage implementation generates significant improvements in labor force participation and economic productivity, with estimated effects ranging from 8.3% to 15.7% across the seven developing countries in our analysis [44]. These impacts emerge gradually over three to five year periods as populations adjust to expanded healthcare access, emphasizing the importance of dynamic modeling approaches in policy evaluation. The identification of preventive care access as a primary mechanism suggests that comprehensive primary care coverage proves more effective for generating labor market benefits than approaches focused primarily on catastrophic care insurance.

The substantial heterogeneity in treatment effects across different population groups highlights the potential for targeted implementation strategies to enhance policy effectiveness while managing fiscal costs. Women, older workers, and individuals with chronic health conditions demonstrate systematically larger treatment effects, while geographic variation suggests that areas with limited baseline healthcare infrastructure experience greater labor market improvements. These patterns provide empirical guidance for optimizing universal health coverage design and implementation strategies. [45]

The welfare analysis incorporating labor market benefits alongside direct health outcomes reveals that economic returns to universal health coverage substantially exceed implementation costs across all countries examined, with benefit-cost ratios ranging from 2.1 to 4.6 over five-year evaluation periods. These findings suggest that universal health coverage represents not only a social policy intervention but also an effective economic development strategy that generates sustained productivity improvements and fiscal benefits.

The methodological contributions of this research extend beyond the specific application to universal health coverage evaluation, demonstrating how adversarial training techniques can address selection bias and confounding in observational policy evaluation. The integration of multiple neural network architectures to capture different aspects of causal relationships provides a template for applying deep learning methods to other policy evaluation contexts where traditional methods face identification challenges.

While limitations regarding temporal scope, geographic coverage, and interpretability constrain the generalizability of our findings, the results provide compelling evidence for the potential of machine learning approaches to advance causal inference in policy evaluation [46]. Future research incorporating longer time series, broader country coverage, and enhanced interpretability techniques could further strengthen these methodological contributions while expanding their policy applications.

The policy implications of our analysis suggest that developing countries should prioritize universal health coverage implementation as part of broader economic development strategies, with particular attention to sequential rollout approaches that begin with underserved areas and emphasize comprehensive primary care services. The demonstrated economic returns provide strong justification for international development cooperation focused on healthcare system strengthening, while the evidence for gender equity benefits suggests additional rationale for prioritizing universal health coverage in countries with pronounced gender gaps in economic participation.

The successful application of deep learning methods to this complex policy evaluation challenge demonstrates the maturation of machine learning techniques for causal inference applications. As data availability and computational resources continue to expand, these methodological approaches offer promising avenues for generating more precise and policy-relevant evidence on the impacts of social policy interventions. The integration of domain expertise from health economics with cutting-edge machine learning techniques represents a productive direction for advancing both methodological development and policy-relevant research in development economics. [47]

References

[1] W. Qiu, C. Chu, A. Mao, and J. Wu, "The impacts on health, society, and economy of sars and h7n9 outbreaks in china: A case comparison study," *Journal of environmental and public health*, vol. 2018, pp. 2710185–2710185, 6 2018.

[2] S. P. Wallace, "The relevance of a political economy of aging analysis in a social determinants of health era," *Innovation in Aging*, vol. 2, pp. 369–369, 11 2018.

[3] Y. Yang, "Does economic growth induce smoking?—evidence from china," *Empirical Economics*, vol. 63, no. 2, pp. 821–845, 2022.

[4] N. Păun and E. Buvilski, "Decision making in central africa- the public health system in light of the changing economy," *On-line Journal Modelling the New Europe*, pp. 4–17, 12 2019.

[5] V. C. Wong, *The Political Economy of Health Care Development and Reforms in Hong Kong - State Intervention in Health Care*, pp. 31–54. Routledge, 5 2019.

[6] D. Michaels and J. Barab, "The occupational safety and health administration at 50: Protecting workers in a changing economy.," *American journal of public health*, vol. 110, pp. 631–635, 3 2020.

[7] Q. Ran, S.-Y. Lee, D. Zheng, H. Chen, S. Yang, J. Moore, and W. Dong, "How would shifting manufacturing from china to indonesia or india impact human health and social economy?," 5 2023.

[8] U. Bajwa, D. Gastaldo, E. D. Ruggiero, and L. Knorr, "The health of workers in the global gig economy.," *Globalization and health*, vol. 14, pp. 124–124, 12 2018.

[9] "Tehran will prioritise economy over health," 6 2020.

[10] B. Anderson, "A nurse's perspective on the impact of coronavirus on health and the economy.," *British journal of nursing (Mark Allen Publishing)*, vol. 31, pp. 46–49, 1 2022.

[11] M. T. Pham and L. H. Nguyen, "Conceptualizing the bidirectional relationship between macroeconomic stability and public health systems: A systematic analysis of feedback loops and dynamic equilibrium models," *Journal of Big Data Processing, Stream Analytics, and Real-Time Insights*, vol. 15, no. 3, pp. 1–16, 2025.

[12] R. Banerjee, A. A. Maruta, and R. Donato, "Does higher financial inclusion lead to better health outcomes? evidence from developing and transitional economies," *Economics of Transition and Institutional Change*, vol. 31, pp. 363–401, 7 2022.

[13] A. R. Chowdhury, A. Surie, and G. Bhan, "Breastfeeding knowledge and practices of working mothers in the informal economy in new delhi: A formative study to explore new intervention pathways towards improved maternal and child health outcomes.," *Social science & medicine (1982)*, vol. 281, pp. 114070–114070, 5 2021.

[14] M. Davis and E. Hoyt, "A longitudinal study of piece rate and health: evidence and implications for workers in the us gig economy.," *Public health*, vol. 180, pp. 1–9, 12 2019.

[15] J. Chepkorir, N. Agata, N. Kiambi, and B. Nangehe, "Institutionalizing leadership management and governance for health system strengthening in emerging economies: Evidence from the partnership for health system strengthening in africa (phssa)programme," *European Journal of Business and Management Research*, vol. 6, pp. 47–52, 11 2021.

[16] J. Caudeville, *The Routledge Handbook of the Political Economy of the Environment - Diagnostics and policy tools to measure and mitigate environmental health inequalities*, pp. 320–342. Routledge, 9 2021.

[17] N. Zhao, B. Li, H. Li, G. Li, R. Wu, Q. Hong, P. Mperejekumana, S. Liu, Y. Zhou, R. Ahmad, A. M. I. Zayan, C. Pemberton-Pigott, and R. Dong, "The potential co-benefits for health, economy and climate by substituting raw coal with waste cooking oil as a winter heating fuel in rural households of northern china.," *Environmental research*, vol. 194, pp. 110683–, 1 2021.

[18] N. Liu and S. Chen, "The protection mechanism of personal health information in the digital economy environment.," *Journal of environmental and public health*, vol. 2022, pp. 2314468–, 6 2022.

[19] F. Rodriguez, A. Sarraju, and M. P. Turakhia, "The gig economy worker-a new social determinant of health?," *JAMA cardiology*, vol. 7, pp. 125–125, 2 2022.

[20] A. T. Mhazo and C. C. Maponga, "The political economy of health financing reforms in zimbabwe: a scoping review.," *International journal for equity in health*, vol. 21, pp. 42–, 3 2022.

[21] R. K. Biswas, S. Huq, and A. Afiaz, "Relaxed lockdown in bangladesh during covid-19: Should economy outweigh health?," *International journal of health policy and management*, vol. 9, pp. 488–490, 11 2020.

[22] G. Shi, D. Wang, and M. Altuntaş, "Financial development and health outcomes: Do financial globalization matter in selected asian economies?," *Frontiers in public health*, vol. 10, pp. 843935–, 4 2022.

[23] S. White, B. Warner, P. Mairs, L. Cope, K. Lohoar, W. Melia, G. Sisson, C. Kothari, and S. Birdi, "P394 benefits of a vbic to both patients and the local health economy," *Posters*, vol. 70, pp. A244.1–A244, 1 2021.

[24] A. J. Caban-Martinez, K. M. Santiago, P. L. Feliciano, K. Ogunsina, H. E. Kling, K. Griffin, and N. S. Solle, "Acute musculoskeletal pain reported among rideshare drivers in the health/safety investigation among non-standard workers in the gig economy (h.i.n.g.e.) pilot study.," *Journal of occupational and environmental medicine*, vol. 62, pp. e236–e239, 3 2020.

[25] I. Inna, Y. Boiko, K. Iryna, R. Serhii, S. Ihor, N. Tubaltseva, H. Natalia, and I. Olena, "Development of socialization of the economy on the conditions of financing sports and health activities," *International journal of health sciences*, pp. 2819–2826, 6 2022.

[26] A. Freni-Sterrantino and V. Salerno, "A plea for the need to investigate the health effects of gig-economy.," *Frontiers in public health*, vol. 9, pp. 638767–, 2 2021.

[27] F. Taghizadeh-Hesary and F. Taghizadeh-Hesary, "The impacts of air pollution on health and economy in southeast asia," *Energies*, vol. 13, pp. 1812–, 4 2020.

[28] D. Shipton, S. Sarica, N. Craig, G. McCartney, S. V. Katikireddi, G. Roy, P. McGregor, and G. Scobie, "Knowing the goal: an inclusive economy that can address the public health challenges of our time.," *Journal of epidemiology and community health*, vol. 75, pp. 1129–1132, 6 2021.

[29] I. V. Bukhtiyarov, L. P. Kuzmina, N. P. Golovkova, A. G. Chebotarev, L. M. Leskina, N. S. Khelkovsky-Sergeev, and N. I. Kotova, "Justification of the platform of standards based on the risk's assessment to health employees disorders of the leading sector's enterprises of the economy," *Russian Journal of Occupational Health and Industrial Ecology*, vol. 61, pp. 155–160, 4 2021.

[30] R. Adhikari, *Migrant Health Professionals and the Global Labour Market - The international migration market economy*, pp. 52–74. Routledge, 10 2019.

[31] E. J. Techera and G. Winter, *Marine Extremes : Ocean Safety, Marine Health and the Blue Economy*. Routledge, 1 2019.

[32] P. A. Banks, “Wingfield, a. h. (2019). flatlining: Race, work, and health care in the new economy:,” *Work and Occupations*, vol. 48, pp. 99–101, 10 2020.

[33] E. Ferlie, S. Dopson, C. Bennett, M. Fischer, J. Ledger, and G. McGivern, *Oxford Scholarship Online - The political economy of English public services reform and implications for management knowledges in health care organizations*. Oxford University Press, 11 2018.

[34] K. Croke, M. B. M. Yusoff, Z. Abdullah, A. N. M. Hanafiah, K. Mokhtaruddin, E. S. Ramli, N. F. Borhan, Y. Almodovar-Diaz, R. Atun, and A. K. Virk, “The political economy of health financing reform in malaysia.,” *Health policy and planning*, vol. 34, pp. 732–739, 9 2019.

[35] D. S. Pangestu, null Sukono, N. Anggriani, and N. M. Yaacob, “Quantifying the health–economy trade-offs: Mathematical model of covid-19 pandemic dynamics,” *Computation*, vol. 12, pp. 139–139, 7 2024.

[36] R. Rani and K. Gakhar, “An assessment of economic health of brics economies after economic crises 2008,” *Abhigyan*, vol. 36, pp. 31–39, 6 2018.

[37] B. Tsofa, E. Waweru, J. Munywoki, K. Soe, D. C. Rodriguez, and A. D. Koon, “Political economy analysis of sub-national health sector planning and budgeting: A case study of three counties in kenya.,” *PLOS global public health*, vol. 3, pp. e0001401–e0001401, 1 2023.

[38] E. Giannakis, J. Kushta, D. Giannadaki, G. Georgiou, A. Bruggeman, and J. Lelieveld, “Exploring the economy-wide effects of agriculture on air quality and health: Evidence from europe.,” *The Science of the total environment*, vol. 663, pp. 889–900, 1 2019.

[39] J. J. Miti, M. Perkiö, A. Metteri, and S. Atkins, “Author response for "factors associated with willingness to pay for health insurance and pension scheme among informal economy workers in low- and middle-income countries: a systematic review",” 7 2020.

[40] X. Shi, J. Li, F. Wang, H. Dinçer, and S. Yüksel, “A hybrid decision-making approach for the service and financial-based measurement of universal health coverage for the e7 economies.,” *International journal of environmental research and public health*, vol. 16, pp. 3295–, 9 2019.

[41] A. Jakubowska, “Behavioural health factors and limitations to the health of labour force: Analysis of the convergence process of the eu economies,” *EUROPEAN RESEARCH STUDIES JOURNAL*, vol. XXIII, pp. 258–276, 11 2020.

[42] A. Hajj, D. A. Badro, C. A. Selwan, H. Sacre, R. Aoun, and P. Salameh, “Gender differences in mental health outcomes amid the covid-19 pandemic and a collapsing economy: A cross-sectional study,” 12 2020.

[43] A. Gungoraydinoglu, I. Öztekin, and Özde Öztekin, “The impact of covid-19 and its policy responses on local economy and health conditions,” *Journal of Risk and Financial Management*, vol. 14, pp. 233–, 5 2021.

[44] D. Kernick, “Monoclonal antibodies for migraine prevention: hope, hype, and health economy challenge.,” *The British journal of general practice : the journal of the Royal College of General Practitioners*, vol. 70, pp. e523–e524, 7 2020.

[45] Özlem Arzu Azer and İbrahim Bora Oran, *AI Applications in Health Sector: Use of Artificial Intelligence in Covid-19 Crisis and Impacts of Medical Robots on Global Economy*, pp. 203–224. Springer Nature Singapore, 5 2023.

[46] O. C. Buonomo, D. Vinci, G. D. Carolis, M. Pellicciaro, F. Petracca, A. Sadri, C. Buonomo, M. Dauri, and G. Vanni, “Role of breast-conserving surgery on the national health system economy from and to sars-covid-19 era.,” *Frontiers in surgery*, vol. 8, pp. 705174–, 1 2022.

[47] L. Zhou, S. Ampon-Wireko, H. A. Antwi, X. Xu, M. Salman, M. O. Antwi, and T. M. N. Afua, “An empirical study on the determinants of health care expenses in emerging economies.,” *BMC health services research*, vol. 20, pp. 774–, 8 2020.