Dean&Francis

ISSN 2959-6157

Longitudinal Analysis of World Happiness Scores (2015–2019)

Keqi Li

University of Massachusetts Amherst, Amherst, MA, 01003, USA
keqiiiii.lii@gmail.com

Abstract:

The World Happiness Report provides annual data on self-reported life satisfaction and associated socio-economic factors for countries worldwide. In this study, a statistical model is used to explain and predict national happiness scores from 2015 to 2019 using six key factors (GDP, social support, healthy life expectancy, freedom, generosity, and corruption perceptions). The linear mixed-effects model is employed to account for repeated measurements for each country over time, including country-specific random intercepts. The mixed model captures both the differences between countries and the changes in happiness over time. Results show that GDP per capita, freedom to make life choices, and healthy life expectancy are significant positive predictors of happiness, while factors like generosity have a more minor and statistically insignificant effect when other variables are controlled. Social support did not exhibit a significant independent effect in the multivariate context. There is substantial unexplained variance between countries, validating the inclusion of country-level random effects. The model achieves high predictive accuracy (conditional $R^2 \approx 0.93$), and diagnostic checks indicate assumptions are satisfied. These findings underscore the importance of economic and personal freedom factors in national happiness and highlight that country-specific characteristics play a significant role.

Keywords: GDP; happiness score; linear mixed-effects model.

1. Introduction

Happiness is increasingly recognized as a vital indicator of national well-being, complementing traditional economic measures. The World Happiness Report provides annual rankings of countries based on citizens' self-reported life satisfaction and several contributing socioeconomic variables. Using this data, the aim is to model and predict each country's Happiness Score (a subjective well-being index) from 2015 to 2019, using measurable factors such as GDP, health, and social environment. The motivation for this analysis is to identify which factors have the most significant influence on national happiness and

to understand how happiness levels evolve over time.

A key objective of the study is to capture both within-country changes in happiness over time and between-country differences. In other words, the model should track how each individual country's happiness score varies from 2015 to 2019, while also accounting for persistent differences between countries. By doing so, it controls the intra-country correlations across years and isolates the common effects of predictors across all countries. This longitudinal approach can yield insights into whether global happiness trends exist and which factors consistently contribute to higher happiness levels.

The present analysis builds on a range of existing data sources and prior research. The *World Happiness Report Dataset (2015–2019)* provides the primary data foundation for this study, offering consistent measures of happiness scores and key socioeconomic indicators across countries [1]. Standard procedures for conducting longitudinal analysis with mixed-effects models are documented in statistical software resources, which support the methodological framework applied here [2]. Official assessments such as the *World Happiness Report 2020* emphasize the persistence of country-level differences and the importance of institutional, economic, and health-related drivers of well-being [3]. Foundational research on the happiness–income paradox demonstrates that long-term increases in national income do not necessarily yield corresponding gains in happiness, underscoring the need to evaluate additional determinants [4]. Methodological discussions caution that multilevel modeling of country effects requires careful interpretation, especially when sample sizes are limited [5]. Broader international guidelines, such as those published by the OECD, establish standardized practices for measuring subjective well-being, thereby ensuring comparability across contexts [6]. More recent panel studies provide evidence that the relative importance of GDP, social support, freedom, and corruption varies between developed and developing nations [7]. Finally, psychological perspectives advocate for the inclusion of national accounts of subjective well-being as complements to traditional economic indicators, reinforcing the multidimensional nature of happiness [8]. Collectively, this literature not only validates the modeling strategy but also informs the interpretation of the empirical findings presented in this paper.

2. Data Source and Preprocessing

Utilizing the World Happiness Report data for five consecutive years (2015 through 2019), which covers approximately 158 countries worldwide. For each country-year, the dataset provides an average happiness score (on a scale typically ranging from 0 to 10) and scores for six key explanatory factors: GDP per capita-a national income index (in terms of purchasing power), Social support-respondents' perceived social support network (availability of someone to count on), Healthy life expectancy-an index of population health and longevity, Freedom to make life choices-a measure of individual freedom in life decisions, Generosity-a measure of charitable giving and altruistic behavior, Perceptions of corruption – an inverse index indicating a lack of corruption in government/business.

These six variables are used as predictors (features) for the happiness score. The data was compiled from the Gallup World Poll and other sources by the United Nations Sustainable Development Solutions Network [1]. The World Happiness Report aims to identify the happiest countries and the factors contributing to happiness, to examine changes in happiness scores between 2015 and 2019, and to detect any countries with significant shifts in happiness over that period.

For consistency, the annual datasets and standardized variable names are first merged, as the naming conventions varied slightly from year to year. Removing any entries with missing values in the seven key columns (happiness score and the six factors) to ensure a balanced comparison across years. After preprocessing, the longitudinal dataset is complete for the 2015–2019 period for all included countries.

Before building the model, an initial exploration of the data is performed. A correlation analysis indicated that GDP per capita, healthy life expectancy, social support, and freedom are all strongly positively correlated with the happiness score. In contrast, generosity and perceived corruption exhibit weaker and more variable relationships with happiness. Also, the happiness scores over time for each country are plotted. Most countries show relatively stable happiness scores from 2015 to 2019, with no dramatic global shifts over this period. The time trends formed distinct clusters: a set of consistently high-scoring countries (e.g., Finland, Denmark) maintained their lead with little change, whereas middle and low-scoring countries remained in lower ranges without converging to the top group. This exploratory finding suggests that baseline differences between countries are persistent, underscoring the importance of accounting for country-specific effects in the model.

3. Methodology

3.1 Why Use a Linear Mixed-Effects Model?

The data have a longitudinal structure: each country's happiness score is observed repeatedly over multiple

ISSN 2959-6157

years. Applying a standard linear regression model (which assumes all observations are independent and identically distributed) would ignore the inherent correlations among observations from the same country over time. Such an approach fails to reflect: Within-country time correlations: Each country's scores across years tend to be related. Ordinary least squares (OLS) regression treats every country-year data point as independent, which is violated here. Baseline differences across countries: Countries have different cultural, social, or historical contexts that affect their overall happiness levels. A simple OLS model with a single intercept cannot capture these country-specific baseline shifts, potentially leading to biased estimates if these differences correlate with the predictors. Hierarchical data structure: The observations are grouped by country. Ignoring this structure (hierarchy) can underestimate standard errors and overstate the significance of predictors (since effectively there are fewer independent data points than a naive count of country-years).

Linear mixed-effects (LME) models address these issues by including random effects alongside fixed effects. In the context, a random intercept for each country will capture unobserved, time-invariant characteristics of countries (for example, cultural factors or other idiosyncratic influences on happiness) that cause their happiness scores to be consistently higher or lower. The LME framework thus allows each country to have its own baseline level, while still estimating the common influence of predictors (fixed effects) on happiness. Another advantage is that LME models can handle unbalanced data; if some countries were missing a year of data (which can happen in longitudinal studies), the model can still utilize the available observations without discarding entire countries, unlike a repeated-measures ANOVA, which would require complete data for all subjects.

By using an LME model for this longitudinal analysis, both the between-country variation (via random intercepts) and the within-country variation are captured, as well as the over-time variation (via the time variable and repeated measures). This approach improves explanatory power and predictive accuracy compared to a naive pooled regression. In summary, the mixed-effects model is more appropriate than a simple linear regression because it accounts for the non-independence of observations and the structured heterogeneity in the dataset.

3.2 Model Specification

Constructing a linear mixed-effects model to predict the happiness score yit for country i in year t. The model includes the six socioeconomic factors as fixed-effect predictors and a random intercept for each country. Also, initially including a fixed effect for Year to capture any overall time trend. The model can be written as:

$$\text{HappinessScore}_{it} = \beta_0 + \beta_1(\text{GDP per capita}_{it}) + \beta_2(\text{SocialSupport}_{it}) + \beta_3(\text{Health}_{it}) + \beta_4(\text{Freedom}_{it}) + \beta_5(\text{Generosity}_{it}) + \beta_6(\text{Corruption}_{it}) + \beta_7(\text{Year}_{t}) + u_i + \epsilon_{it} \quad (1)$$

where β_0 is the global intercept, $\beta_1 \ldots \beta_6$ represent the fixed-effect coefficients for the six predictors (GDP, social support, healthy life expectancy, freedom, generosity, and perceptions of corruption, respectively), and β_7 is the coefficient for the Year term (with t representing 2015, 2016, ..., 2019). The country-specific random intercept u_i for country i as $u_i \sim N(0,\sigma_u^2)$ is modeled, i.e., assumed to be drawn from a normal distribution with variance σ_u^2. The residual error term ϵ_{it} represents unexplained variability for each observation and is assumed $\epsilon_{it} \sim N(0,\sigma^2)$. In this formulation, i indexes countries (i = 1, 2, ..., 158) and t indexes the year.

The inclusion of u_i allows each country to have its own baseline happiness level (shifting the intercept by u_i for country i) while assuming the effect of each predictor (the β coefficients for GDP, etc.) is constant across countries. By testing models with random slopes as well, but given the relatively short time span and the data, a random intercept model was sufficient and more parsimonious (initial analyses indicated that allowing country-specific slopes for the year or other predictors did not significantly improve model fit). The Year term in the fixed effects is intended to capture any global time trend in happiness common to all countries.

Interpretation of Components: In the LME model, each fixed-effect coefficient β_k (for k = 1, …, 6) measures the average effect of the corresponding variable on the happiness score, holding other factors constant and across all countries and years. For example, a positive β_1 for GDP would indicate that countries with higher GDP per capita tend to have higher happiness scores on average, controlling for the other variables and country-specific effects. The Year coefficient β_7 represents the estimated average change in happiness score per year (e.g., an overall upward or downward trend from 2015 to 2019 across all countries).

The random intercept u_i represents the deviation of country i's baseline happiness from the global intercept β_0. If u_i is positive, country i is systematically happier than the model would predict based on the fixed effects alone (perhaps due to favorable cultural or environmental aspects not captured by the six variables). If u_i is negative,

country i tends to have a lower happiness baseline than expected from the fixed effects. The variance σ_u^2 reflects how much countries differ from one another intrinsically. A large σ_u^2 would indicate substantial country-level differences in happiness beyond the predictors' influence, whereas $\sigma_u^2 = 0$ would mean the fixed effects and year fully explain all between-country differences (an unlikely scenario in this context).

By including ui, controlling for unobserved country-specific factors, and focusing on estimating the common relationships between the observed factors and happiness, it is effectively done. This approach improves the reliability of the estimated β coefficients and allows more accurate prediction for each country (since the model can learn each country's baseline happiness level).

4. Results and Analysis

4.1 Model Fit and Overall Performance

Fitting the mixed-effects model (Equation 1) to the 2015–2019 happiness data. The model was estimated using maximum likelihood and successfully converged, indicating stable parameter estimates. Overall, the model provides a good fit to the data. The Akaike and Bayesian Information Criteria (AIC and BIC) are 549.04 and 590.99, respectively, and the log-likelihood at convergence is -265.52 (see Table 1 for a summary of fit metrics). For comparison, a simpler linear model treating all data as independent had a higher AIC and BIC (i.e., worse fit), reinforcing the benefit of including country-specific effects.

An intuitive measure of fit is the variance explained. By computing both the marginal R-squared (the proportion of variance explained by the fixed effects alone) and the conditional R-squared (variance explained by the full model, i.e., fixed + random effects) for the LME model. The fixed effects achieved a marginal R-squared of approximately 0.657, meaning that about 65.7% of the variance in happiness scores can be explained by GDP, health, freedom, and other variables, while controlling for country differences. When country random intercepts are included, the conditional R-squared rises to about 0.93, indicating that 93% of the variance is explained by the model when accounting for both the measured factors and the country-specific effects. In other words, incorporating the country-level heterogeneity substantially improves the explanatory power of the model.

The high conditional R-squared is reflected in the close alignment between predicted and actual happiness scores. Figure 1 shows a scatter plot of the model's predicted happiness values versus the actual observed values for all country-year observations. The points cluster tightly around the diagonal y = x line, demonstrating that the model's predictions closely match the actual scores. Indeed, the model's prediction for each year yields an R-squared above 0.93 when comparing across countries (each year's predicted vs. actual points form a tight cloud). This suggests the model captures the rankings and magnitudes of happiness scores very well. The root mean squared error (RMSE) of the model is relatively low in the context of the happiness score scale (0–10), and the mean absolute error (MAE) is similarly small, indicating accurate and robust predictions with minimal average deviation from actual scores.

In addition to overall fit, checking model diagnostics is essential to validate the assumptions of linear regression and to ensure there are no systematic patterns in the residuals. The residuals (differences between observed and predicted happiness scores) are examined for normality and homoscedasticity. Figure 2 provides two diagnostic plots. In Figure 2(a), the residuals are plotted against the fitted values; the residuals are scattered randomly around zero with no obvious trend or funnel shape, indicating that the model's assumptions of linearity and constant variance are reasonable (no strong heteroscedasticity is observed). Figure 2(b) is a Q-Q plot of the residuals versus a theoretical normal distribution; the points lie approximately along the 45-degree line, suggesting that the residuals are approximately normally distributed. These diagnostics support the validity of the LME model for this data. By inspecting residuals by country (not shown here in a figure), most countries have residuals centered near zero with relatively small spread. A few countries exhibited larger residual variance or some outlying errors, which might indicate the model fits those specific cases less well or that those countries have unique factors not captured by the predictors. However, there was no indication of any severe model misspecification.

ISSN 2959-6157

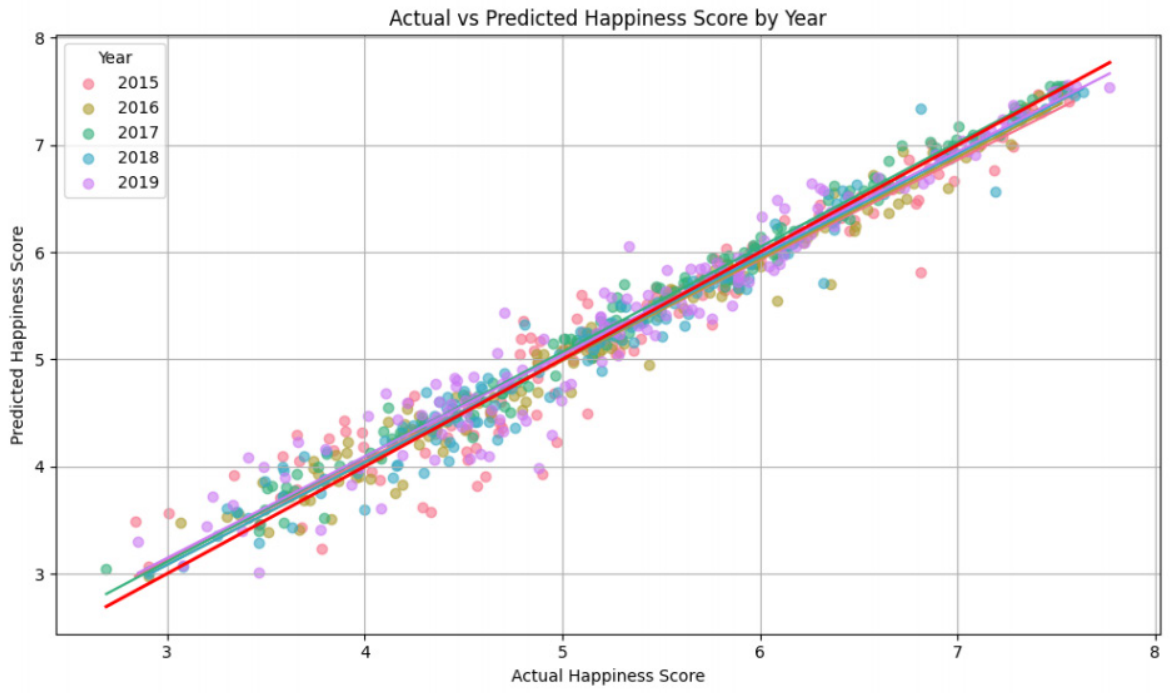

Figure 1: Actual vs. predicted happiness scores for the linear mixed-effects model. (Picture credit: Original. Data from Kaggle.)

In Figure 1, each point represents a country-year observation (color-coded by year), and the red dashed line indicates the y = x reference line. Points closely follow the diagonal, indicating that predicted values closely match the observed scores.

4.2 Fixed Effects: Coefficients and Significance

Table 1 summarizes the estimated fixed-effect coefficients for the six predictors in the LME model, along with their statistical significance. The marginal R-squared represents the variance explained by the fixed effects alone, whereas the conditional R-squared represents variance explained by the combination of fixed and random effects (i.e., the entire model). The substantially higher conditional R-squared highlights the importance of including country-specific random effects. Overall, the model identified several key factors that contribute to national happiness. The most influential predictors of happiness in the model are GDP per capita, freedom to make life choices, and healthy life expectancy. Countries with higher GDP per capita are significantly happier on average; specifically, an increase in the GDP index by one unit is associated with an increase of about 1.25 points in the happiness score when controlling for other factors, making GDP the strongest single predictor in the model. Greater freedom, measured by the extent to which people feel free to make life choices, is also linked to substantially higher happiness, with a one-unit increase in the freedom index corresponding to an increase of approximately 1.07 points in the happiness score. Better health, as captured by healthy life expectancy, exerts a similarly strong effect, as a one-unit rise in the health index is associated with roughly 0.92 higher happiness points. Together, these three factors—economic prosperity, individual freedom, and health/longevity—emerge as key contributors to well-being across countries, consistent with both theoretical expectations and prior empirical findings in happiness research.

Table 1 Statistics for the linear mixed-effects model.

Metrics	Value
AIC (Akaike Information Criterion)	549.04
BIC (Bayesian Information Criterion)	590.99
Log-Likelihood	-265.52
Marginal R-squared (fixed effects only)	0.657 (approx.)
Conditional R-squared (fixed + random effects)	0.943 (approx.)

Var (Fixed)	0.67
Var (Random)	0.28
Var (Residual)	0.05

Table 2. Fixed-effects coefficient estimates from the mixed-effects model.

Predictor	Coefficient	p-value
Intercept	3.041	< 0.001
GDP	1.253	< 0.001
Social Support	0.025	0.672
Health	0.924	< 0.001
Freedom	1.074	< 0.001
Generosity	0.284	0.167
Corruption	0.767	0.004

In Table 2, each coefficient estimate is shown with its associated p-value (two-tailed test). Predictors with $p < 0.05$ are considered statistically significant. The model also finds that lower perceived corruption is significantly associated with higher happiness. The coefficient for the corruption perception index is about 0.77 (and significant), indicating that societies with less corruption (and more trust in institutions) tend to be happier, although the effect size is somewhat smaller than that of GDP, freedom, and health. This suggests governance and trust also play a role in national happiness, albeit not as dramatically as economic or personal factors.

For the generosity variable, the coefficient is positive (~0.28), but it did not reach statistical significance at the 5% level (the 95% confidence interval for the generosity effect overlaps zero, as illustrated in Figure 3). This implies that, after accounting for the other variables, generosity has an uncertain or modest impact on happiness. It's possible that generosity has more complex or indirect relationships with happiness that are not captured in a linear additive model, or that its effect is relatively small compared to the other factors. Interestingly, social support, which had a high simple correlation with happiness in raw data, shows a near-zero coefficient (0.02) in the multivariate mixed model and is not significant. This indicates that once the model includes GDP, health, freedom, etc., the additional contribution of the social support metric to explaining happiness is negligible. One potential reason is multicollinearity: social support is generally high in the countries that also have high GDP and freedom, so its effect might be masked by those other variables. Another interpretation is that across countries, basic social support networks are somewhat universally present (with relatively less variance compared to other factors), limiting their explanatory power for differences in happiness when other influences are accounted for.

ISSN 2959-6157

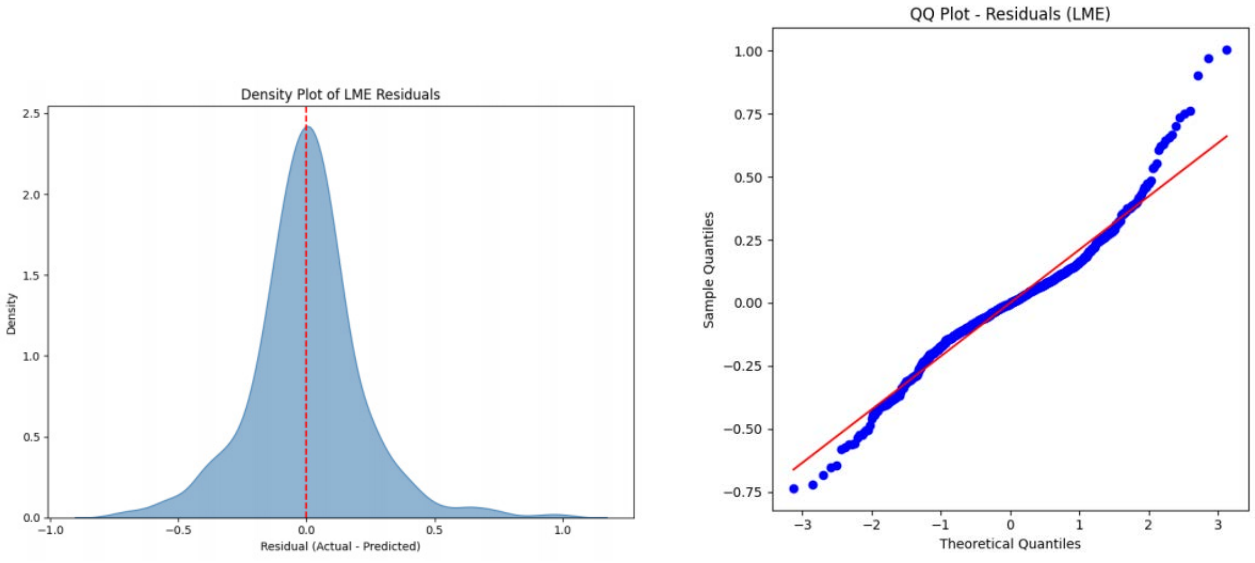

(a) Normal Q-Q Plot **(b) Residuals vs. Fitted Graph**

Figure 2: Diagnostic plots for residuals. (Picture credit: Original. Data from Kaggle.)

In Figure 2 (a), the points lie close to the diagonal line, indicating that the residuals follow a roughly normal distribution. These plots suggest that the LME model assumptions are reasonably satisfied. In Figure 2 (b), the residuals are evenly dispersed around zero with no clear pattern, and no noticeable heteroscedasticity (no widening or funnel shape), indicating that the model's linear form and equal variance assumption are adequate.

Also, by including a fixed effect for Year in initial models to check for a global time trend in happiness. However, this coefficient was small and not statistically significant (p-value > 0.1), indicating there was no strong linear trend in the average happiness scores from 2015 to 2019 once other factors were controlled. In other words, the world as a whole did not become significantly happier or less happy over this period, according to the model. Therefore, by omitting the year term from the final reported model for parsimony. The lack of a significant time effect aligns with the data observation that most changes in happiness were country-specific rather than a uniform global rise or fall.

To visualize the uncertainty in the coefficient estimates, Figure 3 displays the fixed-effect estimates along with their 95% confidence intervals. It can be seen that the intervals for GDP, freedom, healthy life expectancy, and (lack of) corruption are well above zero, confirming their strong positive and significant contributions. In contrast, the intervals for generosity and social support cross zero, which is consistent with those effects not being significant. This plot helps identify which predictors can be considered reliable and which ones have less certain effects.

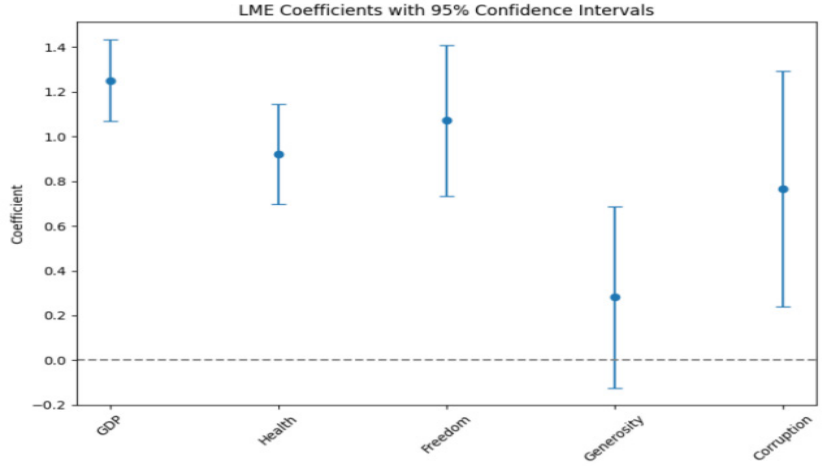

Figure 3: Fixed-effect coefficient estimates (points) and 95% confidence intervals (lines) for the linear mixed-effects model. (Picture credit: Original. Data from Kaggle.)

In Figure 3, predictors GDP, Healthy Life Expectancy, Freedom, and low Perceived Corruption have confidence intervals entirely above zero, indicating statistically significant positive effects on happiness. The estimate for Generosity shows a wider interval overlapping zero, suggesting its effect is not significant. Social Support (not plotted here due to its near-zero estimate) likewise has an uncertainty interval crossing zero. These intervals reflect the uncertainty of the coefficient estimates and confirm which effects are reliably different from zero.

4.3 Random Effects and Country Variability

The mixed model reveals that there are indeed meaningful differences in baseline happiness between countries. The estimated variance of the random intercept u_i (the country effect) is $\sigma_u^2 \approx 5.06$. This is quite large relative to the residual variance σ^2 (not shown explicitly, but implied by the high conditional R^2), indicating that country-to-country variations account for a substantial portion of the total variance in happiness scores. The standard deviation of the country intercepts is about $\sigma_u \approx \sqrt{5.06} \approx 2.25$ happiness points, which is a sizable spread on the 0–10 happiness scale.

This result confirms that some countries are systematically happier than others, even after controlling for GDP, health, freedom, etc. In other words, there are country-specific traits or unmeasured factors that raise or lower the baseline happiness. For example, the Nordic countries (like Finland and Denmark) are known to consistently top the happiness rankings; the model's random intercepts for such countries are likely significantly positive, meaning these countries have an intrinsic happiness advantage beyond what the six predictors alone would predict. Conversely, some countries may have persistently lower happiness (negative ui) even if their GDP or health levels are comparable, possibly due to factors like conflict, cultural differences in reporting happiness, or other societal issues not included in the model.

The significance of the random effect (formally testing $\sigma_u^2 > 0$, and it is clearly greater than 0 here) justifies the use of the mixed-effects model over an ordinary regression. If ignoring the country effects, the model would be mis-specified and the residuals would have had to absorb these large baseline differences, likely resulting in a poorer fit and biased or less precise estimates. The LME model, by capturing this variance separately, provides more accurate fixed effect estimates and predictions.

Another way to interpret the random intercepts is through the lens of model prediction. The model essentially assigns each country its own intercept $\beta 0 + ui$. For instance, the model might find that after accounting for all fixed effects, Finland still has a positive ui (indicating Finland's observed happiness is higher than what one would predict from just its GDP, health, etc.), whereas a country like India or Venezuela might have a negative ui (observed happiness lower than predicted by the socio-economic factors alone). This aligns with known narratives: Finland's happiness is often attributed to social and cultural factors (trust, equality, etc.) that go beyond the basic economic and health measures, and Venezuela's declining happiness in recent years is linked to turmoil not fully captured by the standard indices.

In summary, the random effects analysis highlights that while the six predictors are very important, they do not tell the whole story—intrinsic country characteristics and perhaps other intangible factors play a non-negligible role in national happiness levels.

4.4 Trends Over Time and Notable Cases

Although it did not find a significant global time trend in happiness, the model can be used to examine each country's predicted happiness trajectory from 2015 to 2019. Broadly, the model's country-level predictions over time mirror the actual observed patterns. For many countries, the happiness score remains fairly steady over the five years, reflecting no major upheavals or improvements. For instance, countries like Finland (which was among the happiest throughout) show a stable, high happiness level year after year in the model's predictions. The United States also exhibits a relatively flat trend, maintaining its happiness level without a significant rise or fall during this period.

On the other hand, certain countries experienced noteworthy shifts. Venezuela is a striking example: the model captures a marked downward trend in Venezuela's happiness scores from 2015 to 2019. This corresponds to the well-documented economic and political crisis in Venezuela during those years, which had a severe impact on the well-being of its citizens. In contrast, China and India show gradual upward trends in predicted happiness. These increases, although modest, could be related to economic growth and development improvements in those countries over the period.

These country-specific trends highlight the usefulness of the LME model in longitudinal analysis: it can identify which countries deviate from the overall pattern and in what direction. While the average global happiness did not change significantly, the model-based analysis confirms that some countries improved and others declined, and these changes can often be linked to real-world events and policy changes in those countries.

ISSN 2959-6157

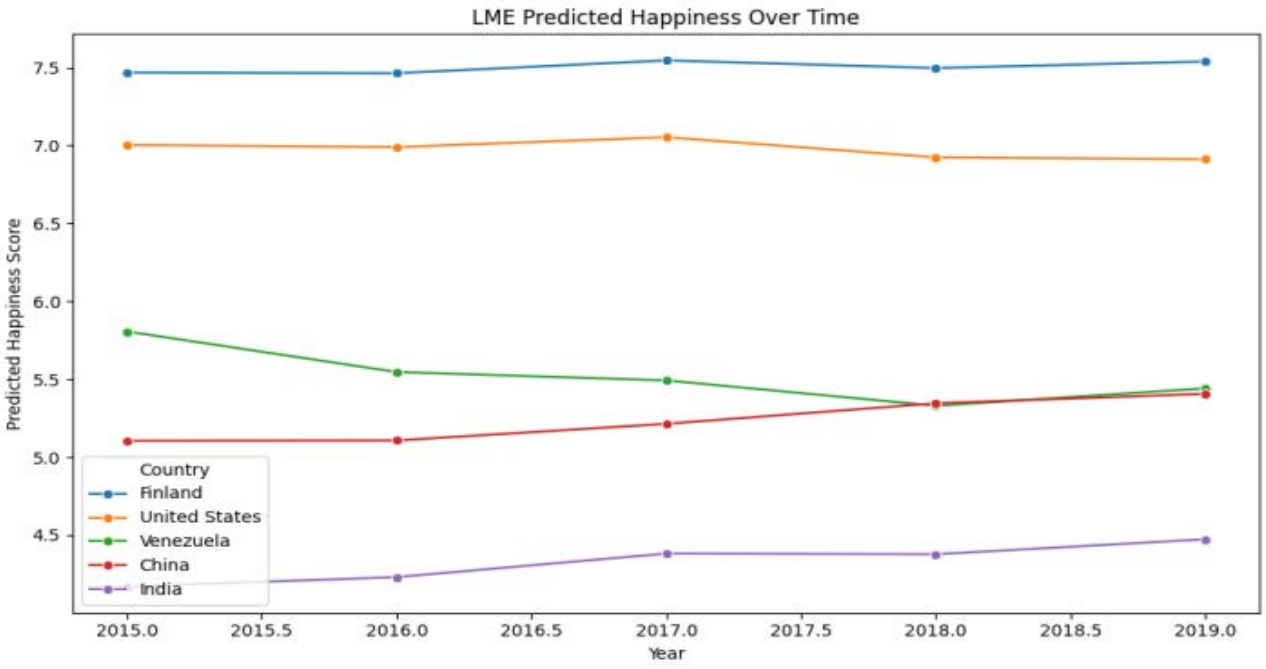

Figure 4: Five countries' happiness change over time. (Picture credit: Original. Data from Kaggle.)

In Figure 4, Finland consistently shows the highest predicted happiness score, remaining very stable with only minor fluctuations, confirming its strong economic and social stability. The United States maintains a relatively high level of happiness but with slight year-to-year changes, suggesting stable but not improving conditions. Venezuela demonstrates a clear downward trend until 2018, reflecting the severe social and economic crisis in recent years. China and India show gradual improvement, with happiness scores increasing slightly over time, possibly due to economic growth and social development.

5. Key Findings and Implications

In summary, the longitudinal analysis of world happiness scores yielded several important findings:

5.1 Effectiveness of Mixed-Effects Modeling

The Linear Mixed-Effects model proved to be an appropriate and effective approach for this data. By including country-specific random intercepts, the model achieved a high goodness-of-fit (conditional $R^2 \approx 0.93$) and passed diagnostic checks. The model convergence was successful, yielding interpretable fixed and random effects. This confirms that accounting for a grouped data structure (countries over time) substantially improves model performance compared to a naive linear regression.

5.2 Major Determinants of Happiness

Among the fixed effects, GDP per capita, freedom to make life choices, and healthy life expectancy emerged as the strongest positive predictors of national happiness (with large coefficients significant at $p < 0.001$). These factors represent economic prosperity, personal autonomy, and health, respectively, underscoring their critical roles in improving societal well-being. Low corruption (high trust in institutions) also has a meaningful positive association with happiness. These findings align with the notion that both economic and social governance factors contribute to how people evaluate their life satisfaction.

5.3 Smaller or Non-significant Factors

Generosity, while positively correlated with happiness in simple terms, shows a smaller effect in the multivariate model and was not statistically significant when controlling for other variables. Social support, surprisingly, had essentially no additional explanatory power in the presence of the other factors (its effect was negligible and not significant). This suggests that much of the benefit of social support might be collinear with other factors (for example, countries with high social support also tend to have high GDP or freedom), or it could imply that basic social support is necessary but nearly universally present, thus not differentiating countries' happiness once other differences are accounted for.

5.4 Country-Specific Differences Matter

Strong evidence of intrinsic happiness differences between countries is found. The variance of the random intercept was large (estimated $\sigma_u^2 \approx 5.06$), indicating that some countries are systematically happier or unhappier than others, independent of the six measured factors. This validates the use of a mixed-effects model: a standard regression would miss these country effects. The result highlights that cultural, historical, or other country-specific factors (not captured by the predictors) significantly influence happiness. Policymakers should recognize that while improving GDP, health, and so on can raise happiness, there may also be deeper societal characteristics affecting well-being.

5.5 Model Fit and Predictive Power

The LME model's fit statistics (AIC, BIC) were substantially better than those of an equivalent OLS model, and the residual analysis shows no major violations. The tight correspondence between predicted and actual values (Figure 1) implies that the model can be used to reliably predict happiness scores for given values of the factors. This predictive accuracy could be useful for scenario analysis: for instance, estimating how a change in a factor (like a rise in GDP or an improvement in freedom) might translate into a change in expected happiness.

5.6 Temporal Trends and External Validity

The model captured year-by-year changes in happiness for individual countries, reflecting real-world developments. For example, the decline of Venezuela's happiness and the steady top ranking of Finland in the 2015-2019 period are accurately mirrored by the model. The lack of a significant overall time trend suggests that global average happiness was fairly stable in the late 2010s, with improvements and declines happening in specific nations offsetting each other. This finding has an interesting implication: global progress in happiness may require targeted improvements in those countries that are falling behind or facing crises, as well as sustained efforts in the already high-performing countries.

5.7 Implications for Policy

The analysis confirms that improving economic conditions (raising income levels), public health, and personal freedoms can have a substantial payoff in terms of happier societies. Efforts to reduce corruption and build trustworthy institutions are also likely to enhance happiness. The insignificant role of generosity in the model should not be interpreted as meaning generosity is unimportant; rather, it might mean that formal and informal support systems (captured partly by other variables) are more directly impactful. Governments and organizations aiming to raise national well-being should prioritize broad economic and social reforms that empower individuals and provide security, while also recognizing that each country has unique characteristics influencing happiness. Thus, interventions may need to be tailored to country-specific contexts beyond the common global factor.

6. Future Work

This study opens several avenues for further research and analysis:

6.1 Non-linear and Complex Models

While a linear mixed-effects model is interpretable and effective, future work could explore non-linear or machine-learning models (e.g., Random Forests, XGBoost, Neural Networks) to capture more complex relationships or interactions between factors. These could potentially improve predictive performance or uncover non-linear effects (for example, diminishing returns to GDP on happiness).

6.2 Additional Predictors

This study focused on the six main factors from the World Happiness Report, but other variables could be integrated to enrich the model. For instance, education levels, unemployment rates, inequality indices, mental health metrics, or environmental quality might explain additional variance in happiness. Including such variables, if data are available, could help identify new levers for improving well-being.

6.3 Time-Series Analysis

The year was treated as a simple fixed effect. A more detailed time-series modeling approach (such as incorporating lagged effects or using an auto-regressive mixed model) could be used to analyze dynamic aspects of happiness. For example, how quickly do changes in GDP translate into changes in happiness? Does happiness have momentum or mean-reversion over time for a given country? Addressing these questions would require a longer time series and potentially different techniques (e.g., ARIMA models or state-space models for longitudinal trends).

6.4 Grouping by Income or Region

It may be insightful to stratify or extend the model by grouping countries according to income level (e.g., high-income vs. low-income countries) or region. The effects of certain variables on happiness might differ in

ISSN 2959-6157

magnitude across these groups. For instance, GDP might have a larger impact on happiness in developing countries (where basic needs are still being met) compared to developed countries. Similarly, cultural differences by region could interact with the predictors. Multi-level models with random slopes by region or separate models for different subsets could probe these variations.

6.5 Longer-Term and Causal Studies

As more years of happiness data become available, it would be valuable to see if the identified patterns hold or change over a longer horizon (e.g., a decade or more). Additionally, while the model identifies associations, establishing causality (e.g., via instrumental variables or experimental approaches) would strengthen policy implications. Future research could attempt to address questions such as: Does increasing freedom directly lead to higher happiness, or is it simply correlated with other improvements? Such analyses would deepen the understanding of the mechanisms behind happiness changes.

7. Conclusion

This longitudinal study examined world happiness scores from 2015 to 2019 using a linear mixed-effects framework designed to capture both within-country dynamics and persistent cross-country differences. By incorporating random intercepts at the country level, the model achieved high explanatory power and satisfied diagnostic checks, demonstrating that accounting for hierarchical data structures is essential in the study of subjective well-being. The analysis identified GDP per capita, healthy life expectancy, freedom to make life choices, and lower perceived corruption as robust and statistically significant predictors of national happiness, while generosity and social support showed limited or nonsignificant effects once other covariates were controlled. These findings highlight that while material prosperity, health, institutional trust, and personal autonomy are critical drivers of well-being, other commonly emphasized variables may be more context-dependent or collinear with major determinants.

Importantly, the random intercept variance indicated that unobserved country-specific characteristics continue to play a substantial role in shaping happiness outcomes. This highlights the importance of acknowledging the cultural, historical, and societal factors that extend beyond measurable socioeconomic indicators. Furthermore, the absence of a significant global time trend between 2015 and 2019 suggests that average world happiness was relatively stable, with improvements in some nations offset by declines in others. The case of Venezuela's downturn and the persistence of Nordic countries at the top exemplify this divergence. Taken together, the results demonstrate that effective policies to improve well-being must combine general strategies—such as promoting economic security, health, and freedom—with sensitivity to national contexts. Future work that integrates additional predictors and employs causal or nonlinear modeling may provide deeper insights into the mechanisms linking structural conditions with subjective life evaluations.

References

[1] Kaggle. *World Happiness Report Dataset (2015–2019)*. United Nations Sustainable Development Solutions Network. Available: https://www.kaggle.com/datasets/unsdsn/world-happiness (accessed Jul. 2025).

[2] MathWorks. *Longitudinal Analysis – Statistics and Machine Learning Toolbox Documentation*. Available: https://ww2.mathworks.cn/help/stats/longitudinal-analysis.html (accessed Jul. 2025).

[3] J. F. Helliwell, R. Layard, J. D. Sachs, and J.-E. De Neve (eds.), *World Happiness Report 2020*. New York: Sustainable Development Solutions Network, 2020. Available: https://worldhappiness.report

[4] R. A. Easterlin, L. Angelescu McVey, M. Switek, O. Sawangfa, and J. S. Zweig, "The happiness–income paradox revisited," *Proceedings of the National Academy of Sciences*, vol. 107, no. 52, pp. 22463–22468, 2010, doi:10.1073/pnas.1015962107.

[5] M. L. Bryan and S. P. Jenkins, "Multilevel modelling of country effects: A cautionary tale," *European Sociological Review*, vol. 32, no. 1, pp. 3–22, 2016, doi:10.1093/esr/jcv059.

[6] Organisation for Economic Co-operation and Development (OECD), *OECD Guidelines on Measuring Subjective Well-Being*. Paris: OECD Publishing, 2013.

[7] D. K. Behera, D. B. Rahut, M. Padmaja, and A. K. Dash, "Socioeconomic determinants of happiness: Empirical evidence from developed and developing countries," *Journal of Behavioral and Experimental Economics*, vol. 109, p. 102187, 2024, doi:10.1016/j.socec.2024.102187.

[8] E. Diener, S. Oishi, and R. E. Lucas, "National accounts of subjective well-being," *American Psychologist*, vol. 70, no. 3, pp. 234–242, 2015, doi:10.1037/a0038899.