# Design, Monitoring, and Impact Evaluation of Chile's Universal Immunization Strategy against Respiratory Syncytial Virus

Leonardo J. Basso[1,2], Felipe del Solar[1,2], Gonzalo Díaz[1,3], Marcel Goic[1,2], Ignasi Neira[1,2], Miguel O'Ryan[1,4], Denis Sauré[1,2], Charles Thraves[1,2*], Juan Pablo Torres[1,4], Natalia Trigo[1,5], Amal Zgheib[1,2]

[1*]Instituto Sistemas Complejos de Ingeniería (ISCI), Santiago, Chile.
[2]Industrial Engineering Department, Universidad de Chile, Santiago, Chile.
[3]SC Johnson College of Business, Cornell University, Ithaca, NY, USA.
[4]Facultad de Medicina, Universidad de Chile, Santiago, Chile.
[5]Anderson School of Management, University of California at Los Angeles, Los Angeles, CA, USA.

*Corresponding author(s). E-mail(s): cthraves@uchile.cl;

**Abstract**

Respiratory Syncytial Virus (RSV) is the leading cause of severe lower respiratory tract infection in infants and saturates pediatric intensive-care capacity in Chile every winter. In 2023 nirsevimab, a single-dose monoclonal antibody, made universal protection medically possible but not yet available or affordable: the manufacturer had no plans to enter Chile before 2028, and the Ministry of Health required evidence of net benefit before reallocating funds. We describe the operations research behind the world's first nationwide universal RSV immunization strategy. Because the quantities driving the decision were unobservable, each was constructed through optimization. Integer programming recovered the national RSV burden series that laboratory records do not provide. A Dirichlet-calibrated Monte Carlo simulation propagated the ordered efficacy intervals reported in the phase III nirsevimab clinical trial across 13 candidate strategies, and a break-even sensitivity analysis identified the price and healthcare-cost conditions under which each strategy remained cost-saving. A two-stage integer program matched high-risk infants, among whom 99% coverage left almost no controls. Augmented

synthetic control reconstructed the unobserved 2024 season from 266 donor series. The program reached 94% coverage among 155,000 eligible infants; effectiveness was 76% against hospitalization and 85% against intensive care unit (ICU) admission; and no infant died of RSV in 2024 or 2025, a first in Chile's recorded history. The counterfactual analysis attributes to the program 59,000 fewer hospital bed-days, 22,500 fewer ICU bed-days, and a net benefit of US$20 million. The strategy improved equity, relieved peak ICU pressure, and saved money. The framework has been replicated in Paraguay and is informing strategies in Colombia and Peru.

**Keywords:** OR in health services, Integer programming, Monte Carlo simulation, Augmented synthetic control, Immunization strategy design

# 1 Introduction

Respiratory Syncytial Virus (RSV) is the leading cause of severe lower respiratory tract infection (LRTI) in infants and a primary driver of pediatric hospitalization worldwide (Li et al., 2022; Shi et al., 2017). In Chile the burden is sharply seasonal. Every winter RSV saturates pediatric intensive care unit (ICU) capacity and forces the health system into a temporary expansion of beds and staff: in 2023, occupancy of pediatric ICU beds for respiratory causes peaked at 888 against a nominal national capacity of 520. Until 2023 the only available prophylaxis was palivizumab, an antibody so costly, and requiring so many doses, that between 2019 and 2023 it reached 12,579 infants — about 1.3% of births.

**The decision.** Nirsevimab, a long-acting monoclonal antibody conferring season-long protection from a single dose, changed what was possible. It did not change what was affordable. In mid-2023 Chile's Ministry of Health (MoH) confronted a decision with two obstacles, both of them evidentiary. The manufacturer had not scheduled Chilean market entry before 2028 and would reprioritize production only for a committed national buyer. And buying at national scale meant reallocating roughly US$45M inside a fixed budget, which the Ministry would not do without a credible estimate of what the program would return. Neither the commitment nor the estimate could come first; each waited on the other. What broke the deadlock was an analysis, delivered in November 2023, showing that in seasons resembling the recent severe ones a universal strategy would have paid for itself on direct medical costs alone — and that in any season it would have removed the winter capacity crisis that no budget line could solve.

**Why the analysis was hard.** Every quantity the decision turned on was unobservable in the data available. Chile confirms RSV by laboratory testing only in sentinel hospitals, so the national series of RSV hospitalizations — the very burden the program was meant to reduce — is not recorded anywhere. The savings from a program never run in 2019, 2022, or 2023 had to be reconstructed from seasons in which it did not exist. After the rollout, measuring effectiveness in the high-risk infants who had previously received palivizumab required comparable non-immunized infants, of

whom the program had left very few — coverage among high-risk infants in the seasonal cohort reached 99%. And valuing the program at the population level required the 2024 RSV season that did not happen.

The recurring task, in other words, was to construct what the data did not show. Standard epidemiological practice constructs such quantities by assumption: fixed code lists drawn from the literature, deterministic point estimates of efficacy, propensity scores, before–after comparisons against historical averages. We instead posed each construction as an optimization problem over national administrative data, and this is the through-line of the paper. The unobserved RSV series became a binary quadratic program over diagnostic codes, fitted to sentinel positivity. The unmeasured joint uncertainty in trial efficacy became a constrained calibration of a Dirichlet distribution that enforces, by construction, the ordering that severity implies. The missing comparison group became a two-stage integer program that maximizes the number of matched cases before minimizing residual imbalance. The unobserved 2024 season was reconstructed using a synthetic control built from a donor pool of 266 unaffected age–disease series.

**Relation to prior work.** Each of the four constructions has a counterpart in the operations research literature on health-care decision making, though not, to our knowledge, assembled into a single decision-analytic pipeline spanning design, approval, implementation, and ex-post evaluation of a national immunization program. Matching under scarce controls has been posed before as an integer or mixed-integer program (Rosenbaum, 1989; Zubizarreta, 2012); our two-stage formulation differs in solving first for the maximum feasible number of matches and only then optimizing balance, rather than the reverse. Resource allocation and capacity planning for infectious-disease control and immunization programs is a long-standing OR application area (Kayvanfar, Baldacci, & Govindan, 2025; Lopes et al., 2022; Mallor & Azcárate, 2014; Samii, Pibernik, Yadav, & Vereecke, 2012); the ex-ante evaluation of Section 3.2 contributes to this stream a break-even sensitivity analysis that reports, rather than a single expected net benefit, the price and cost conditions under which the recommendation would reverse. Synthetic control methods originated in comparative case studies (Abadie, Diamond, & Hainmueller, 2010) and have since been applied to operations and policy evaluation; the donor pool used here (266 series against 274 pre-treatment weeks) is large relative to what the original method assumes, which is why we adopt the augmented variant (Ben-Michael, Feller, & Rothstein, 2021; Doudchenko & Imbens, 2016). Within that variant, we weight the augmented outcome transformations by their XGBoost variable-importance scores (Section 3.5) rather than treating them uniformly as in the original method, which is, to our knowledge, a new extension of ASCM in its own right. What is otherwise new is not any one formulation in isolation but their integration, on a single set of linked national registries, across the full policy cycle.

**Approach and contributions.** A multidisciplinary team of engineers, clinicians, and epidemiologists partnered with the MoH to answer four questions spanning the full policy cycle. We summarize each with the methodological difficulty it posed.

*(i) Which strategy, if any, is worth adopting?* Thirteen candidate strategies, indexed by the size of the catch-up cohort, were compared using national data from 2019–2023. Integer programming selects the International Classification of Diseases, Tenth Revision (ICD-10) code set that best reproduces sentinel RSV positivity. A Dirichlet-calibrated Monte Carlo simulation then propagates through each strategy the wide and logically ordered efficacy intervals reported in the phase III MELODY trial of nirsevimab, rather than relying on the deterministic efficacy point estimates commonly assumed in cost-effectiveness models. Rather than a single expected net benefit, a break-even sensitivity analysis reports the price per dose and the healthcare-cost scaling factor at which each strategy's recommendation would reverse, stating the decision the Ministry actually faced instead of asserting that a projected saving would hold.
*(ii) Did it work at national scale?* Effectiveness against hospitalization and ICU admission was estimated on the full national cohort with stratified Cox proportional hazards models, monitored in real time through a public platform integrating daily data from 234 hospitals, and validated against a negative control outcome and two sensitivity analyses.
*(iii) Did it protect high-risk infants?* These infants are few and their events fewer, which rules out survival models, and near-universal uptake among them leaves almost no controls. A two-stage integer program builds the largest feasible matched sample before optimizing its balance, recovering 173 of 179 cases.
*(iv) What did the country save?* Augmented synthetic control methods (Ben-Michael et al., 2021; Doudchenko & Imbens, 2016) extend the synthetic control approach (Abadie et al., 2010) to a donor pool large relative to the pre-treatment window, yielding counterfactual 2024 series for bed-days, outpatient visits, and medical leave, and hence a valuation of the program in the year it ran.

**What happened.** The Ministry approved the program, secured the doses, and immunized 94% of approximately 155,000 eligible infants between April and September 2024. Nirsevimab prevented 76% of RSV-related hospitalizations and 85% of ICU admissions. For the first time in Chile's recorded history, no infant died of RSV — in 2024 or again in 2025. The counterfactual analysis attributes to the program a reduction of about 59,000 hospital bed-days and 22,500 ICU bed-days and a net economic benefit of US$20M against US$45M in acquisition costs; the Ministry did not have to mount the temporary winter ICU expansion it had run every year for decades. Paraguay replicated the framework for its own 2025 season; Colombia and Peru are preparing strategies for 2026.

The clinical and epidemiological findings have been reported in four medical journals (Saure et al., 2025; Sauré, Torres, et al., 2026; Sauré, Zgheib, et al., 2026; Torres et al., 2025). This paper documents the operations research that produced them: the formulations, the calibration and validation choices, and the places where the methods had to depart from standard practice. Our aim is to make the framework usable by researchers and policymakers facing the same problem elsewhere, which — given a seasonal pathogen, a new prophylactic, and a national health registry — many of them do.

The rest of the paper is organized as follows. Section 2 describes the national data sources. Section 3 presents the four methodological frameworks. Section 4 reports

results. Section 5 details the impact on the Ministry of Health and the broader policy landscape. Section 6 concludes.

## 2 Data

All four analyses run on the same national administrative registries, linked through the unique national identifier assigned at birth and anonymized by the MoH before delivery. The ex-ante analysis of Section 3.2 was performed at the end of 2023 and therefore uses only data available up to that year, excluding immunization records. We describe each source in turn.

**Inpatient care records.** Provided by the MoH, these are collected from the nationwide hospital discharge system. The records encompass all hospital admissions (HA) in both public (156 hospitals) and private (78 hospitals) institutions, and contain patient-related data (e.g., sex, birth date), hospitalization details (e.g., ICD-10 diagnosis codes at admission and during the stay), and data on bed usage including dates of entry/exit and bed type (basic, intermediate, or ICU). Additionally, for 2023 only, the MoH provided a weekly series of ICU bed occupancy, planned capacity, and realized capacity for respiratory causes.

**Outpatient care records.** We retrieved open-access data on daily medically attended (MA) visits to offices and emergency departments (not requiring a hospital admission) at each healthcare facility in the public sector, collected and maintained by the MoH. Data were not available for the private sector, which represents approximately 23.5% of national RSV-related hospital admissions among patients under one year of age. Outpatient data contain information on patient age, initial diagnoses, date of visit, and facility name. Unlike inpatient data, outpatient data are available only through aggregate ICD-10 categories (e.g., "Pneumonia" encompasses codes J12–J18) and two infant age groups: under 1 year and between 1 and 4 years.

**Birth records.** All births in Chile from 2019 to 2024 were provided by the Department of Statistics and Information of the MoH. These records include weeks of gestational age (wga), birth weight, and birth length, facilitating the identification of high-risk groups.

**Immunization records.** These comprise all immunizations for infants, including an anonymized patient identifier, the location, and the date of each administration. We focus on nirsevimab administrations. Because the registries are linked at the individual level, immunization status can be joined to birth characteristics and to subsequent hospital episodes for the entire national birth cohort, rather than for a sampled subset.

**Cost parameters.** Unit costs for emergency department (ED) visits, pediatric basic/intermediate care, and ICU bed-days were provided by the MoH. Although these costs correspond to the public healthcare system, we applied the same values to the private sector, which is a conservative assumption since private network costs are likely higher. Additional cost components include the prices of palivizumab (provided by the MoH, as it had been routinely used for high-risk infants) and nirsevimab (US$225 per dose, the price reported for the Galicia program; Ares-Gómez et al. 2024). We further included an administration cost of US$5 per dose for infants in the catch-up cohort, and zero for those in the seasonal cohort. The former reflects approximately half of the

administration cost reported in Canada; the latter assumes negligible logistical costs when administering nirsevimab to newborns at birth in healthcare facilities.

# 3 Methodological Approaches

The project, named NirseCL, adopted an operations research framework that integrated integer programming, Monte Carlo simulation, Cox proportional hazards modeling, and augmented synthetic control methods with large-scale administrative data to design, implement, and evaluate Chile's national RSV immunization strategy. The work unfolded across several phases, all of which required estimating daily counts of MA, HA, and ICU RSV-associated LRTI (data not directly observed in national health records). We begin by describing how these series were constructed, as they underpin all subsequent analyses; we then address each of the four problems stated in Section 1. Not every stage of this pipeline is itself an optimization problem: Sections 3.1, 3.2, and 3.4 construct or select a decision-relevant quantity by solving one, while Section 3.3 instead closes the loop by measuring, on the realized 2024 season, whether the decision those constructions supported performed as designed. Figure 1 summarizes the resulting four-stage pipeline — design, approval and implementation, verification, and ex-post evaluation — mapping each stage to the method that carries it and to the series-construction problem of Section 3.1 that three of the four analytical stages depend on.

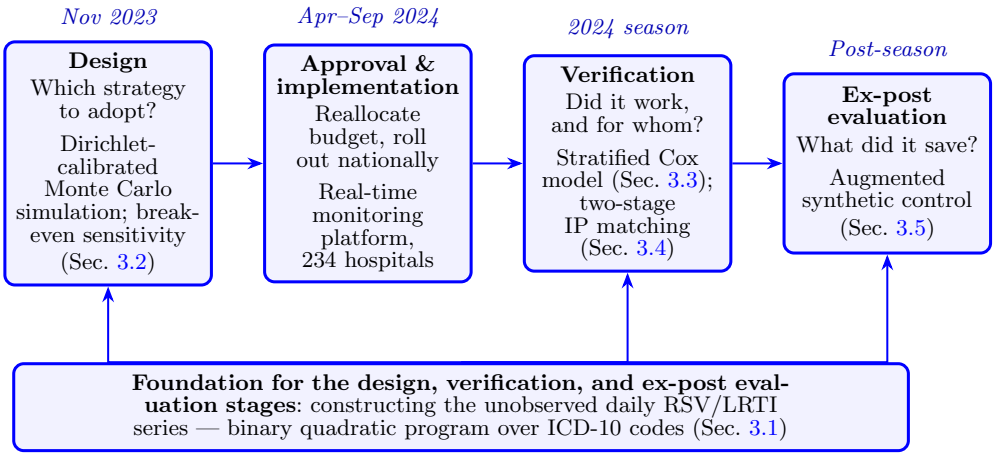


**Fig. 1**: The decision-analytic pipeline linking the four operations research constructions of Section 3 to the policy cycle: design under uncertainty (Sec. 3.2), the Ministry of Health's approval and national rollout, verification against the realized 2024 season (Secs. 3.3–3.4), and ex-post economic evaluation (Sec. 3.5). The code-selection problem (Sec. 3.1) constructs the daily RSV/LRTI series that the design, verification, and evaluation stages all require as input.

### 3.1 Constructing daily series of HA and ICU RSV-associated LRTI

Chile's nationwide sentinel surveillance program comprises 32 hospitals, ten of which test every LRTI admission by immunofluorescence with real-time polymerase chain reaction (PCR) confirmation. Outside those centers, RSV is not routinely confirmed. This is the first and most consequential gap in the data: ICD-10 discharge codes are observed for every admission in the country, but laboratory-confirmed RSV positivity is observed only for a sentinel subset. Every downstream quantity — the burden to be averted, the savings to be claimed, the effectiveness to be measured — rests on how RSV admissions are identified in the other 200-odd hospitals.

To overcome this gap, we estimated daily RSV-related hospital admissions and bed occupancy (by bed type) by selecting the subset of ICD-10 primary discharge codes most likely to correspond to RSV infection. We considered three nested candidate sets, labeled RSV0, RSV1, and RSV2, each progressively expanding the previous set. RSV0 includes codes J12.1, J20.5, and J21.0; RSV1 adds J21.9 and B97.4; RSV2 further adds J12.8, J12.9, J18.0, J18.8, J18.9, J20.8, and J20.9. Construction of these sets was informed by prior studies and clinical expert input. For sentinel hospitals we computed, for each epidemiological week, the proportion of LRTI-related admissions with a primary diagnosis in each candidate set.

To select the best-fitting code set in a principled way, we formulated an integer programming model in which binary decision variables $x_c \in \{0, 1\}$ indicate whether each ICD-10 code $c$ should be classified as RSV-related. The objective minimizes the sum of squared deviations between the sentinel positivity series and the positivity implied by the selected codes:

$$\min_{x} \sum_{t \in T} \left( p_t - \sum_{c \in C} a_{tc}\, x_c \right)^2, \qquad \text{s.t.} \quad x_c \in \{0, 1\} \quad \forall c \in C, \tag{1}$$

where $C$ is the set of candidate codes, $T$ is the set of epidemiological weeks, $p_t$ is the sentinel RSV positivity rate in week $t$, and $a_{tc}$ is the fraction of LRTI-related admissions with code $c$ in week $t$ among sentinel hospitals. This is a binary quadratic program over the full $2^{|C|}$ subsets of the 12 candidate codes, not merely a choice among the three named sets RSV0–RSV2 (which serve as reference points against which the optimized set is compared in Appendix A); for the instance considered here ($|C| = 12$, $|T| = 156$ sentinel weeks) it is small enough to solve to global optimality with standard branch-and-bound methods, in well under a second on a standard laptop using Gurobi. The code set obtained from this optimization coincided exactly with RSV1; we therefore adopted RSV1 as our preferred definition, namely J12.1 (pneumonia due to RSV), J20.5 (acute bronchitis due to RSV), J21.0 (acute bronchiolitis due to RSV), J21.9 (acute bronchiolitis, unspecified), and B97.4 (RSV as the cause of diseases classified elsewhere). An admission is labeled RSV-related when its primary discharge diagnosis lies in this set, it occurs during the RSV season, and it occurs at least 7 days after birth. Figure A1 in Appendix A illustrates the fit of each candidate set against the sentinel positivity series.

Two features of this formulation deserve comment. First, the optimization is not merely a convenience: it converts a judgment call that would otherwise be settled by clinical opinion into a decision with an auditable objective, which mattered when the resulting burden estimates were presented to the Ministry. Second, the selected set includes J21.9, a code that does not name RSV, and its inclusion is the main source of potential misclassification. Restricting attention to codes that explicitly mention RSV would bias effectiveness estimates upward; we therefore retain J21.9 in the primary analysis and report the restricted definition as a sensitivity analysis in Section 4.2.

Using the admissions associated with RSV1, we derived daily bed-occupancy series by type using recorded lengths of stay. A similar methodology was applied to generate daily MA RSV-associated LRTI series, with two key adjustments: private hospital activity was extrapolated (due to absent direct records), and infant ages in months were imputed because MA data are available only in aggregated age groups.

## 3.2 Ex-ante cost-saving evaluation

The cost-saving analysis described here informed the Ministry of Health's decision to implement the immunization program. It evaluates nationwide immunization policies that can simultaneously cover three potentially overlapping groups: the *seasonal cohort*, newborns delivered during the RSV season; the *catch-up cohort*, infants who by March 2024 — that is, before the onset of the season — are younger than a given number of months, that number being the defining parameter of the strategy; and the *high-risk cohort*, infants who meet the criteria for palivizumab.

We use the term *strategy* to refer to the choice of catch-up cohort size. The seasonal and high-risk groups are always included; we vary only the catch-up window, considering 13 options indexed 0–12 by the number of months included. For instance, strategy 6 corresponds to immunizing all infants who will be under 6 months of age by March 2024, in addition to the high-risk and seasonal cohorts. This configuration closely resembles the approach adopted in Galicia, Spain (Ares-Gómez et al., 2024).

**Monte Carlo simulation.** To quantify the direct cost savings that would have been realized had nirsevimab been implemented during the 2019, 2022, and 2023 RSV seasons (2020–2021 are excluded because of the extremely low RSV circulation caused by the COVID-19 pandemic), we conducted a Monte Carlo simulation study with 1,000 replications per strategy. For each replication, a vector of efficacy values against MA, HA, and ICU outcomes is drawn from a calibrated probability distribution (described below). Using the observed daily RSV-associated LRTI series, we identify the cases corresponding to infants covered under the strategy being evaluated. Each identified case is then removed from the series with a probability equal to the sampled efficacy for the relevant outcome. We assumed full coverage, consistent with the Chilean context where immunization coverage among infants under one year of age generally exceeds 95%, and constant protective efficacy through 150 days post-dose, which spans the season for every strategy evaluated. The 95% confidence intervals are constructed by selecting the central 950 simulated values around the mean. All strategies replace palivizumab with nirsevimab in the high-risk cohort, so the avoided palivizumab expenditure enters the comparison as a saving.

**Dirichlet-calibrated efficacy model.** The analysis incorporated efficacy estimates from the phase III MELODY trial (Muller et al., 2023): 76.4% (95% confidence interval [CI]: 62.3–85.2%) against MA, 76.8% (95% CI: 49.4–89.4%) against HA, and 78.6% (95% CI: 48.8–91.0%) against ICU RSV-associated LRTI. Because these estimates carry substantial uncertainty and are logically ordered (efficacy against HA is no lower than against MA, and efficacy against ICU is no lower than against HA), we modeled them jointly through a four-category Dirichlet distribution. Each draw yields four positive components summing to one: the first represents sampled efficacy against MA; the second the incremental protection from MA to HA; the third the increment from HA to ICU; the fourth a residual ensuring the components sum to one. By the aggregation property of the Dirichlet, the marginal distribution of the cumulative sum of the first $i$ components is $\mathrm{Beta}(\bar{\alpha}_i,\, \bar{\alpha}_4 - \bar{\alpha}_i)$, where $\bar{\alpha}_i = \sum_{j=1}^{i} \alpha_j$, so that $\bar{\alpha}_4$ is the Dirichlet concentration parameter. This construction guarantees the nesting constraint by design.

Let $\mu_i$ and $p_i(r)$ denote the reported mean and $r$-th quantile for outcome $i \in \{1, 2, 3\}$ (MA, HA, ICU). The four Dirichlet parameters $\{\alpha_j\}_{j=1}^{4}$ are obtained by solving:

$$\min_{\alpha} \sum_{i=1}^{3} \sum_{r \in \{0.025,\, 0.975\}} \Big( p_i(r) - F^{-1}\big(r;\, \bar{\alpha}_i,\, \bar{\alpha}_4 - \bar{\alpha}_i\big) \Big)^2, \tag{2}$$

subject to $\bar{\alpha}_i = \mu_i \cdot \bar{\alpha}_4$ for $i \in \{1, 2, 3\}$ and $\alpha_j > 0$ for all $j$, where $F^{-1}(\cdot;\, a, b)$ is the inverse cumulative distribution function (CDF) of a $\mathrm{Beta}(a, b)$ distribution. This is a smooth nonlinear program in four variables, solved using standard gradient-based methods. The calibrated parameters and the resulting fit to the MELODY trial confidence intervals are reported in Appendix B.

**Break-even sensitivity.** Two parameters carried most of the decision risk: the price the Ministry would ultimately negotiate per dose, which was not yet known, and the unit healthcare costs, which are public-sector averages applied to a private sector whose true costs are proprietary. Rather than report a single net figure, we inverted the calculation in each direction. For fixed healthcare costs we computed the maximum price per dose at which a strategy remains cost-saving; for a fixed price per dose we computed the multiplicative factor by which healthcare costs would have to be scaled to make it so. Both quantities are reported in Section 4.1, and both proved more persuasive in discussion with the Ministry than the point estimates, because they state the conditions under which the conclusion fails rather than asserting that it holds.

**Peak capacity.** In addition to cost savings, we used the 2023 nationwide daily series of available and occupied ICU beds due to LRTI to quantify how each immunization strategy would have reduced peak pressure on critical-care capacity during the winter surge. This last quantity proved decisive: the binding constraint in 2023 was not the cost of care but the physical capacity of the pediatric ICU network, which no budget line could expand on short notice.

### 3.3 Effectiveness of nirsevimab

Having designed the strategy in Section 3.2 and matched capacity to it operationally, the remaining task in the design-to-evaluation pipeline was to verify, on the realized season, that the decision performed as the ex-ante model predicted; this section and Section 3.4 report that verification, and Section 3.5 its economic counterpart. The national RSV immunization program ran from April 1 to September 30, 2024. To ensure transparency and continuous monitoring, a public website was developed (NirseCL Team, 2024) integrating real-time data from 234 public and private hospitals nationwide. Once the season concluded, we estimated the effectiveness of nirsevimab using a stratified Cox proportional hazards model (Torres et al., 2025) for two outcomes: (i) RSV-related LRTI hospitalization and (ii) RSV-related LRTI hospitalization requiring ICU admission.

For each outcome, eligible infants in the seasonal cohort were followed from birth until the earliest of: RSV-related disease, death, or September 30, 2024. Infants in the catch-up cohort were followed from April 1, 2024, under the same censoring rules; those who experienced the outcome before this date were excluded. To prevent misclassification of infections occurring before full protection, infants who developed RSV within 7 days after immunization were classified as non-immunized.

We adopted a stratified Cox proportional hazards model in which the stratification variables (sex, gestational age in weeks, age in months, rurality, and region) are identified as potential confounders via a directed acyclic graph. These variables are incorporated through stratification rather than as covariates because they violated the proportional-hazards assumption (Schoenfeld residual tests). The hazard function for infant $i$ is:

$$\lambda\big(t \mid \mathcal{S}_i\big) = \lambda_0\big(t; \mathcal{S}_i\big) \cdot \exp\big(\beta_{\text{imm}} \cdot X_{\text{imm},i}\big), \tag{3}$$

where $\mathcal{S}_i$ indexes the stratum, $\lambda_0(t; \mathcal{S}_i)$ is the stratum-specific baseline hazard, $X_{\text{imm},i}$ is a binary indicator of immunization status, and $\beta_{\text{imm}}$ is estimated jointly across all strata. Effectiveness is then $1 - \exp(\hat{\beta}_{\text{imm}})$.

Subgroup analyses were carried out by fitting separate stratified Cox models for sex, cohort (seasonal or catch-up), prematurity, rurality, and macrozone. Operational differences in healthcare utilization were assessed by comparing hospital length of stay between immunized and non-immunized infants using the Mann–Whitney $U$ test.

**Validation.** Observational effectiveness estimates of this kind are vulnerable to a specific failure: if immunized and non-immunized infants differ in their propensity to be hospitalized at all, the model attributes that difference to the antibody. The concern is not hypothetical here, since non-immunized infants were on average older, heavier, more likely to be born to foreign mothers, and less likely to live in the central macrozone. We therefore refit the identical model on a negative control outcome — hospital admission for intestinal infectious disease (ICD-10 A08.0–A08.5, A09.0, A09.9), which nirsevimab cannot plausibly prevent. An estimate indistinguishable from zero indicates that the stratification absorbs the differential propensity to be admitted; a large estimate would indicate that it does not. We report two further robustness checks: a restricted definition of the outcome that excludes the non-specific code J21.9, and an imputation of the 2,289 infants whose immunization status could not be determined and who are excluded from the primary analysis.

### 3.4 Effectiveness in high-risk infants

Infants were classified as *high-risk* if they met either of the following conditions: (i) extreme prematurity, defined as gestational age $< 32$ weeks or birth weight $< 1{,}500$ g; or (ii) congenital heart disease (CHD), identified by the presence of a hospital discharge record with a CHD ICD-10 code within the infant's first year of life. Infants born preterm but not meeting the high-risk criteria (gestational age $< 36$ weeks) were classified as *non–high-risk preterm*. The union of both groups defines the *at-risk* population.

Real-world evidence on nirsevimab effectiveness in high-risk infants remained limited to clinical trials at the time of this analysis (Griffin et al., 2020; Hammitt et al., 2022; Simões et al., 2023). This subgroup is the one the previous policy existed to protect, so the question of whether a universal strategy served it at least as well was the crux of the policy change — and the hardest to answer. High-risk infants are few, their events fewer, which rules out the stratified survival models of Section 3.3; and coverage among them reached 99%, so the non-immunized comparators that any design needs had nearly ceased to exist. We therefore adopted a matched case–control design with conditional logistic regression (Sauré, Torres, et al., 2026), in which the scarcity of controls becomes an allocation problem: each control can be used once, and the design must decide which cases to spend them on.

**Two-stage optimization matching.** The matching was performed with a two-stage integer programming strategy designed to (i) maximize the number of eligible case–control matches on blocking variables, then (ii) minimize residual covariate imbalance, while enforcing strict comparability and no control reuse.

Let $I$ and $J$ denote the sets of at-risk infants in the case and control groups, $A_i \subseteq J$ the controls compatible with case $i$ on a set of categorical blocking variables on which exact agreement is required, $B_j = \{i \in I : j \in A_i\}$ the cases compatible with control $j$, and $E = \{(i, j) : j \in A_i\}$ the set of feasible arcs. Let $d_{ij}$ be a standardized distance between case $i$ and control $j$ on the remaining continuous covariates.[1] Targeting a 1:$k$ case-to-control ratio, where $k$ is chosen according to the size of the case and control pools in each analysis, the two integer programs (IPs) are:

[1] The blocking variables align cases and controls on determinants of access to and uptake of immunization rather than on downstream illness characteristics; the distance covers the remaining covariates that plausibly affect baseline susceptibility. The concrete variables used in each analysis, and the covariates carried into the conditional logistic regression, are reported in Sauré, Torres, et al. (2026).

$$
\begin{aligned}
M^{*} = \max_{x,y} \quad & \sum_{(i,j)\in E} x_{ij} \\
\text{s.t.} \quad & \sum_{j\in A_i} x_{ij} = k\, y_i \quad \forall i \in I \\
& \sum_{i\in B_j} x_{ij} \le 1 \quad \forall j \in J \\
& x_{ij},\, y_i \in \{0,1\}
\end{aligned}
\tag{4}
$$

$$
\begin{aligned}
\min_{x,y} \quad & \sum_{(i,j)\in E} d_{ij}\, x_{ij} \\
\text{s.t.} \quad & \sum_{j\in A_i} x_{ij} = k\, y_i \quad \forall i \in I \\
& \sum_{i\in B_j} x_{ij} \le 1 \quad \forall j \in J \\
& \sum_{(i,j)\in E} x_{ij} \ge M^{*} \\
& x_{ij},\, y_i \in \{0,1\}
\end{aligned}
\tag{5}
$$

Decision variable $x_{ij}$ equals 1 if case $i$ is paired with control $j$; $y_i$ equals 1 if case $i$ is paired with exactly $k$ controls. The first IP maximizes the number of matched cases; the second minimizes total covariate distance subject to achieving $M^{*}$ matches.

The lexicographic ordering is the point of the construction. Greedy or nearest-neighbor matching, the default in this literature, processes cases sequentially and can exhaust the controls compatible with a rare stratum before reaching the cases that need them, discarding cases that a global allocation would have retained. Solving for $M^{*}$ first guarantees that no case is lost to the order in which cases happen to be considered, and optimizing balance second ensures that this guarantee costs nothing in comparability beyond what the feasible set already imposes. In a cohort this sparse the difference is material: the procedure retained 173 of 179 cases, and the six it dropped were infeasible rather than crowded out. Both integer programs are solved to global optimality in well under a second on a standard laptop using Gurobi, reflecting the modest size of the instances involved (179 cases and up to 1,211 feasible case–control arcs). We note in passing that, because each case draws exactly $k$ controls and each control is used at most once, the constraint structure of both IPs is that of a bipartite transportation problem and hence totally unimodular; the linear relaxation is therefore already integral, and formulating (4)–(5) as an integer program is a matter of convenience rather than necessity.

With the matched sample, a conditional logistic regression was applied controlling for immunization status (assessed 7 days before the event), sex, wga, and birthweight. Effectiveness was expressed as $(1-\text{odds ratio})\times 100$, with 95% confidence intervals. As a sensitivity analysis, controls were restricted to infants hospitalized for a predefined set of unrelated diagnoses, so that cases and controls are equated on having reached hospital care at all (Carbajal et al., 2024).

### 3.5 Ex-post economic impact evaluation

The effectiveness estimates of Sections 3.3 and 3.4 answer whether nirsevimab protected the infants who received it. They do not answer what the country saved, which is the question the Ministry had to defend at the next budget cycle, and which requires the 2024 season that did not occur.

**Why a simple counterfactual was not enough.** Our first estimate of program impact used a ratio estimator available within weeks of the season's end: for each historical season we computed the ratio of RSV hospitalizations among infants in their first RSV season to those in their second, applied the historical ratio to the observed 2024 second-season count, and read off the first-season cases that did not occur (Torres et al., 2025). The estimator is transparent and was useful for early communication, but it rests on the assumption that the first-to-second-season ratio is stable across years, it uses a single comparison group, it inherits whatever year-to-year variation in RSV circulation the historical ratios happen to contain, and it extends to no outcome other than hospital admissions. It cannot value bed-days, outpatient visits, or medical leave, and it offers no principled uncertainty quantification. We therefore turned to a method that constructs the comparison series rather than assuming it.

**Augmented synthetic control.** We estimated counterfactual outcomes for the 2024 winter season (Sauré, Zgheib, et al., 2026) using the Augmented Synthetic Control Method (ASCM; Ben-Michael et al. 2021; Doudchenko and Imbens 2016). ASCM extends the synthetic control approach of Abadie et al. (2010) by allowing negative weights and adding a ridge-regression bias correction that is important when the donor pool is large relative to the pre-treatment period.

We constructed counterfactuals for four outcome series in infants under one year of age: (i) days of medical leave (ML), (ii) outpatient visits (MA), (iii) basic/intermediate bed-days (BM), and (iv) ICU bed-days. Because RSV is under-coded outside sentinel hospitals, outcomes here are defined over the broader category of LRTI and related viral agents (ICD-10 B95–B98, J09–J18, J20–J22) rather than the narrow RSV1 set of Section 3.1: a definition that undercounts RSV would attribute part of the program's effect to the residual category and bias the estimated impact downward. For each outcome, the donor pool consists of age–disease groups unlikely to have been affected by the immunization campaign. To avoid spillover effects, children aged 1–4 years with LRTI were excluded from the donor pool, which comprised 266 time series. Dummy donor units absorb month and holiday effects, and all series are normalized over the pre-intervention period so that sources on very different scales remain comparable. The estimation window covers 274 pre-intervention weeks (January 2019 to March 2024); the post-intervention window spans 35 weeks (April to November 2024). Although the campaign opened on April 1, the first recorded dose was administered on March 25, so the intervention date is set at March 31, 2024 to preclude anticipation effects.

ASCM proceeds in two steps. First, it finds the convex combination of donor series minimizing pre-treatment prediction error (the synthetic control method (SCM) weights $\mathbf{w}^{\text{SCM}}$). Then it solves:

$$\mathbf{w}^{\text{AUG}} = \arg \min_{\mathbf{w}:\, \sum w_i = 1} \left\{ \frac{1}{2\lambda} \sum_{t=1}^{T_0} \left( y_{1t} - \sum_{i=2}^{N} y_{it} w_i \right)^2 + \frac{1}{2} \sum_{i=2}^{N} \left( w_i - w_i^{\text{SCM}} \right)^2 \right\}, \qquad (6)$$

where $y_{1t}$ is the treated series, $y_{it}$ the $i$-th donor series in week $t$, $T_0$ the last pre-intervention week, and $\lambda$ a regularization parameter selected via cross-validation over placebo tests in the pre-intervention period ($\lambda = 10$ for all outcomes). The donor series are augmented with $K$ covariates and $M$ transformations of the outcome (e.g.,

cumulative monthly counts), with transformation weights derived from XGBoost variable-importance coefficients— an extension of ASCM to weighted feature augmentation that, to our knowledge, is new in its own right, motivated by a donor pool (266 series) large enough relative to the pre-treatment window ($T_0 = 274$ weeks) that treating all augmented features as equally informative would waste degrees of freedom the ridge correction of Equation (6) cannot recover; see Appendix C for details.

Model fit is assessed using the mean error (ME) between synthetic and observed values in the pre-intervention period. The treatment effect is the cumulative difference between the synthetic and observed series over the 35 post-intervention weeks. Confidence intervals are obtained from the distribution of placebo residuals; $p$-values are derived via placebo tests (Bruhn et al., 2017). Monte Carlo simulations decompose effects by cohort (catch-up vs. seasonal).

**Validation.** A synthetic control that fits the pre-treatment period well can still fit it for the wrong reasons, and the resulting counterfactual is then an artifact of the fitting procedure rather than an estimate of what would have happened. We used three checks. First, placebo tests in time: the intervention date was reassigned to the 2022 and 2023 winter seasons, where no effect should be found, both to calibrate the null distribution used for inference and to confirm that the method does not manufacture effects out of ordinary seasonal variation; the same exercise was used to select $\lambda$ and to choose ASCM over several competing variants. Second, inspection of the donor weights: the method is free to select any of the 266 series, so we asked whether the series it selects are epidemiologically sensible ones. Third, comparison of pre- and post-intervention ME: if the post-intervention divergence is of the same order as the pre-intervention residual, the estimated effect is indistinguishable from misfit. Section 4.4 reports all three.

# 4 Results

## 4.1 Ex-ante cost-saving evaluation

Figure 2 presents the simulated reductions in hospital bed-days, ICU bed-days, MA, and total savings for the different immunization strategies, assuming RSV activity levels observed in 2019, 2022, and 2023. The results exhibit an increasing but concave pattern: as more infants are immunized, reductions grow but at a diminishing rate, consistent with the decreasing RSV risk of older infants.

Both the intensity of RSV circulation and the associated pressure on healthcare resources vary considerably by year and resource type. The 2019 season exhibited the highest burden in terms of hospital bed use, whereas 2023 showed the greatest ICU utilization. This distinction is operationally important because ICU bed-days carry a much higher unit cost than basic or intermediate care.

Considering direct cost savings across both public and private healthcare, we estimated that an immunization strategy would have been cost-saving for catch-up cohorts including infants up to 7 months of age in 2023, up to 6 months in 2019, and up to 4 months in 2022. When restricting to savings in the public sector only, the cost-saving thresholds decrease substantially: up to 3 months for 2023, 2 months for 2019, and 1 month for 2022.

Under the joint public–private perspective, strategy 6 would have generated mean net savings of US$2.5M in 2023 (US$47.6M in direct savings vs. US$45.1M in acquisition and administration costs) and US$0.7M in 2019 (US$54.6M vs. US$53.9M). In contrast, it would have resulted in additional costs of US$2.6M in 2022 (US$44.6M in savings vs. US$47.2M in costs).

These margins are thin, which is precisely why the break-even quantities mattered more than the point estimates. Strategy 6 remains cost-saving as long as the negotiated price per dose stays below US$237.67 in a season resembling 2023, US$227.87 in one resembling 2019, and US$212.06 in one resembling 2022 — against a reference price of US$225. Equivalently, healthcare unit costs would have to be scaled by 0.93 and 0.98 for 2023 and 2019 respectively, factors below one, meaning the conclusion survives even if our cost estimates are somewhat generous; for 2022 the required factor is 1.08, so in a mild season the strategy pays for itself only if true costs exceed our estimates by 8%. Stating the result this way made explicit what the Ministry was actually deciding: not whether a projected saving would materialize, but whether it was willing to bear a modest loss in a mild season in exchange for a substantial gain in a severe one — and, in either case, to eliminate the capacity crisis described next.

Figure 3 illustrates the 2023 pediatric ICU bed occupancy against fixed and extended capacity, together with counterfactual occupancy under strategies 0, 6, and 12. ICU occupancy peaked at 888 beds on June 20; one week earlier the extended (winter-campaign) capacity was 858 beds, representing an expansion of 338 beds above the nominal nationwide capacity of 520 pediatric ICU beds. Under strategy 0, occupancy would have peaked at 673 beds (95% CI: 656–693) on June 25; under strategies 6 and 12, peaks were 612 (95% CI: 592–633) and 589 beds (95% CI: 567–614), respectively, both on July 4. These findings were instrumental in convincing all parties to proceed with a nationwide immunization strategy.

## 4.2 Effectiveness of nirsevimab

The national immunization strategy targeted approximately 155,000 eligible infants during the 6-month 2024 RSV season and achieved coverage exceeding 94%: 91.2% in the catch-up cohort and 97.1% in the seasonal cohort. By September 30, 2024, 145,087 infants had been immunized (72,246 in the catch-up cohort, 72,841 in the seasonal cohort). Of the 9,086 non-immunized infants, 6,937 (76.4%) were in the catch-up cohort and 2,149 (23.6%) in the seasonal cohort.

By the end of the RSV season, 957 hospitalizations occurred among 72,128 immunized person-years, compared with 327 hospitalizations over 5,787 person-years among non-immunized infants.

The estimated effectiveness of nirsevimab in preventing RSV-related LRTI hospital admissions was 76.4% (95% CI: 72.6–79.7); see Table 1. Effectiveness against RSV-related ICU admissions was higher, at 84.9% (95% CI: 79.5–89.0), indicating significantly greater protection against severe outcomes — a pattern consistent with the absence of RSV infant mortality in 2024, compared with 13 deaths in 2023.

Stratified analyses revealed significantly higher effectiveness against LRTI hospitalization among females compared with males, and higher effectiveness against

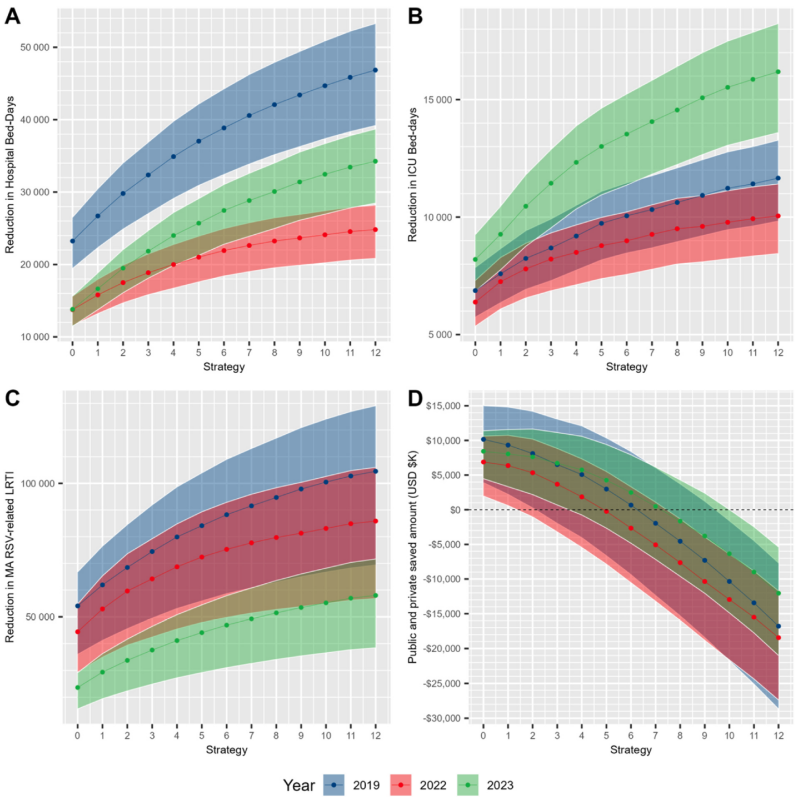


**Fig. 2**: Mean and 95% CI for each immunization strategy (0–12 months catch-up cohort): reduction in hospital bed-days (top-left), reduction in ICU bed-days (top-right), reduction in medically attended (MA) RSV-related visits (bottom-left), and total direct public-and-private savings (bottom-right). Results shown for the 2019, 2022, and 2023 RSV seasons.

both hospitalization and ICU admission in the catch-up cohort relative to the seasonal cohort. The catch-up/seasonal difference for LRTI hospitalization persisted only when the analysis was restricted to full-term infants. No statistically significant differences were observed across rurality or macrozone subgroups, though the highest point estimates were consistently observed in the Central macrozone, which contains the majority of Chile's population. Reliable estimates for the Austral macrozone could not be obtained due to its small population.

Kaplan–Meier survival curves for hospitalization and ICU admission begin to diverge at approximately 75 days (Figure 4). A substantial separation is sustained between immunized and non-immunized infants; confidence intervals are wider for ICU admissions, as expected given lower event counts.

The validation analyses support these estimates. On the negative control outcome, the same model attributes to nirsevimab an apparent 10.1% effectiveness against

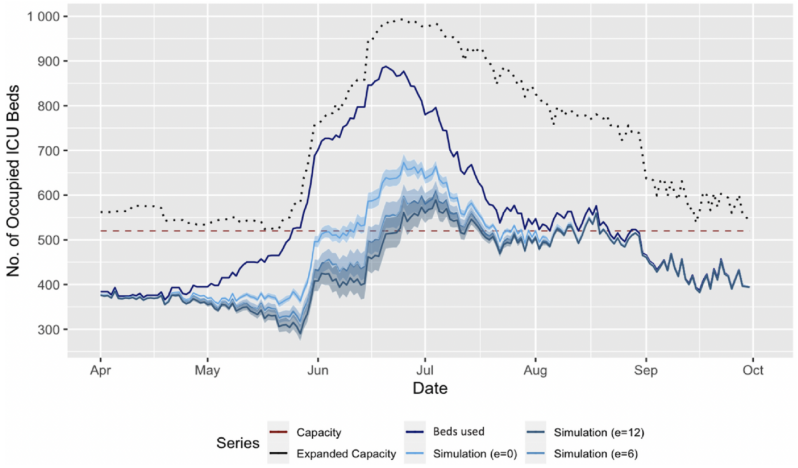


**Fig. 3**: Occupancy vs. expanded capacity of pediatric ICU beds during the 2023 RSV season. The dashed line represents nominal capacity (520 beds); the dotted line represents realized winter-campaign capacity; the dark blue solid line is actual occupancy. Lines with lighter shades represent mean estimated occupancy (and 95% CI) under strategies 0, 6, and 12.

admission for intestinal infectious disease (95% CI: −20.2 to 32.7; $p = 0.5$), statistically indistinguishable from zero: the stratification absorbs the differential propensity to be hospitalized that separates immunized from non-immunized infants, and does not leave a residual that the model could misread as protection. Restricting the outcome definition to codes that explicitly name RSV, and separately imputing the immunization status of the 2,289 excluded infants, leave the point estimates substantially unchanged. Effectiveness also declines monotonically as the outcome becomes less specific to RSV — 76% against RSV-related LRTI hospitalization, 67% against all-cause LRTI hospitalization, and 48% against all-cause hospitalization — which is the gradient a genuine RSV-specific effect should produce, and the opposite of what a generalized healthy-user effect would produce.

Conditional on hospitalization for at least one day, the length of stay differed significantly between immunized (median: 3 days, interquartile range (IQR) 2–5; mean: 4.19 days; $n = 880$) and non-immunized infants (median: 4 days, IQR 2–6; mean: 5.25 days; $n = 312$; Mann–Whitney $U$ test, $p < 0.0001$). For ICU admissions, median stay was 4 days (IQR 2–5) among 181 immunized infants and 4 days (IQR 3–7) among 79 non-immunized infants, with corresponding means of 4.59 and 6.99 days ($p = 0.0062$).

Figure 5 shows cumulative RSV-related LRTI hospital admissions among eligible and non-eligible populations during the 2019, 2022, 2023, and 2024 RSV seasons. By week 25 of 2024, both hospitalizations and ICU admissions among eligible infants are substantially lower than those observed in the non-eligible group, a contrast markedly stronger than in the comparison years.

**Table 1**: Estimated nirsevimab effectiveness against RSV-related LRTI hospital admission and against RSV-related LRTI ICU admission, overall and by subgroup, from the stratified Cox proportional hazards model. Effectiveness $= 1 - \exp(\hat{\beta}_{\text{imm}})$. "Events" is the number of RSV-related LRTI hospital admissions in the subgroup. NA: no reliable estimate could be obtained given the number of events.

| Subgroup | | Events | Effectiveness, % (95% CI) | |
|---|---|---|---|---|
| | | | LRTI admission | ICU admission |
| **Overall** | | 1,284 | 76 (73–80) | 85 (80–89) |
| **Sex** | Female | 483 | 81 (76–85) | 90 (83–94) |
| | Male | 801 | 73 (67–78) | 81 (71–87) |
| **Cohort** | Catch-up | 707 | 81 (77–84) | 88 (82–92) |
| | Seasonal | 577 | 53 (32–68) | 74 (49–87) |
| **Preterm** | No (≥37 weeks) | 1,056 | 77 (73–81) | 85 (80–90) |
| | Yes (<37 weeks) | 228 | 69 (50–80) | 80 (49–92) |
| **Urban/Rural** | Urban | 1,239 | 77 (73–80) | 85 (80–89) |
| | Rural | 45 | 73 (37–89) | NA |
| **Macrozone** | North | 61 | 59 (18–80) | NA |
| | Central | 864 | 79 (74–82) | 86 (80–91) |
| | South | 347 | 74 (66–80) | 86 (76–92) |
| | Austral | 12 | NA | NA |

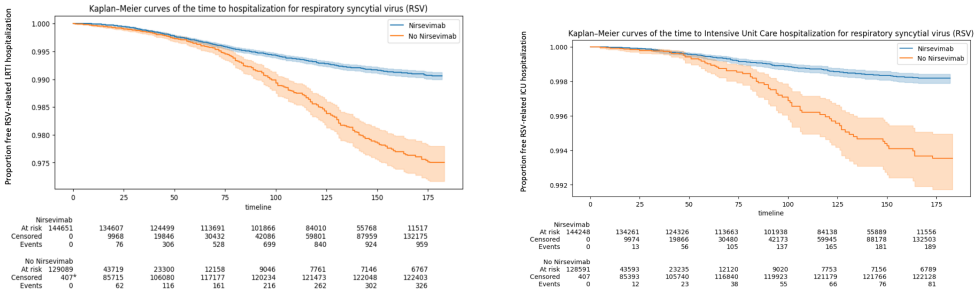


(a) RSV-related LRTI hospitalization.

(b) RSV-related LRTI ICU admission.

**Fig. 4**: Kaplan–Meier curves by immunization status. For immunized infants, follow-up begins at immunization date; for non-immunized infants in the catch-up group, follow-up begins on April 1, 2024; for non-immunized infants in the seasonal group, follow-up begins at birth. Shaded areas represent 95% CIs.

## 4.3 Effectiveness in high-risk infants

The matching procedure included 173 of the 179 at-risk cases, matched to 1,211 controls. For non–high-risk preterm infants, 87 cases were matched to 609 controls; for high-risk infants (extreme preterm or CHD), 86 cases were matched to 602 controls. The six unmatched cases had no feasible control within their blocking stratum, so no

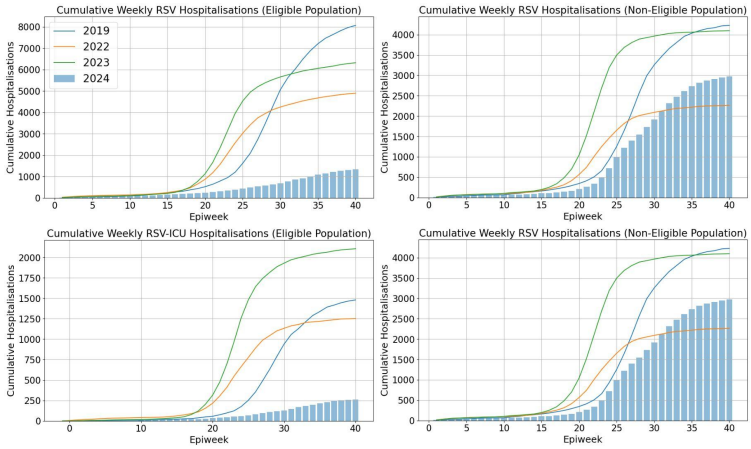


**Fig. 5**: Cumulative RSV-related LRTI hospital admissions (top panels) and ICU admissions (bottom panels) among eligible (left) and non-eligible (right) populations during the 2019, 2022, 2023, and 2024 RSV seasons. The eligible population in each year corresponds to infants that the 2024 strategy would have targeted had it been applied that year. Solid lines depict historical series; bars represent 2024 observed admissions.

allocation of the available controls could have retained them. Baseline characteristics were balanced between cases and matched controls across all groups.

Table 2 reports the results. Overall effectiveness against RSV-related LRTI hospitalization among at-risk infants was 79.5% (95% CI: 62.4–88.8), demonstrating substantial protection even among infants traditionally targeted by palivizumab prophylaxis. In the catch-up cohort, effectiveness among at-risk infants reached 86.8% (95% CI: 73.4–93.5). Reliable estimates could not be obtained for high-risk infants in the seasonal cohort because immunization coverage in this group was 99%, leaving virtually no non-immunized comparators.

Effectiveness for high-risk infants was lower and more uncertain than for the broader at-risk population, reflecting the heterogeneity of this group. Estimates were 87.6% (95% CI: 57.0–96.4) for infants with CHD and 58.6% (95% CI: −18.6 to 85.5) for extremely preterm infants. The lower point estimate for extremely preterm infants may reflect heightened clinical vigilance reducing the threshold for hospital admission when respiratory symptoms are present, irrespective of RSV status. These findings support inclusion of nirsevimab in routine immunization strategies while highlighting extremely preterm infants as a subgroup warranting additional preventive measures or closer clinical monitoring.

## 4.4 Ex-post economic impact evaluation

Figure 6 compares the observed and synthetic series for the four outcomes. Pre-intervention, the synthetic control provides an excellent fit across all outcomes (pre-treatment MEs are small in magnitude; Table 3). Post-intervention, the synthetic

**Table 2**: Nirsevimab effectiveness against RSV-related LRTI hospitalizations in high-risk and at-risk populations. Effectiveness = (1 − odds ratio) × 100. CHD = congenital heart disease; wga = weeks of gestational age.

| Population | | Effectiveness, % (95% CI) | Total cases | Matched cases |
|---|---|---|---|---|
| **At-risk infants** | Overall | 79.5 (62.4–88.8) | 179 | 173 |
| | Catch-up | 86.8 (73.4–93.5) | 112 | 106 |
| **High-risk infants** | Overall | 74.4 (43.6–88.4) | 92 | 86 |
| | Catch-up | 82.4 (55.5–93.1) | 65 | 59 |
| **Extreme preterm** (<32 wga or <1,500 g) | Overall | 58.6 (−18.6 to 85.5) | 58 | 54 |
| | Catch-up | 71.2 (−2.6 to 91.9) | 44 | 40 |
| **Infants with CHD** | Overall | 87.6 (57.0–96.4) | 41 | 35 |
| | Catch-up | 87.0 (48.1–96.7) | 27 | 21 |
| **Non–high-risk preterm** (36 > wga ≥ 32) | Overall | 84.8 (60.1–94.2) | 87 | 87 |
| | Catch-up | 87.9 (66.3–95.7) | 47 | 47 |

series consistently overestimates observed outcomes, particularly for MA, BM, and ICU. The post/pre ME differences are two to three orders of magnitude larger than the pre-treatment residuals against which they must be judged, so the divergence reflects a treatment effect rather than accumulated misfit.

The donor weights are epidemiologically interpretable, which is the check that most directly addresses the concern that a donor pool this large can fit anything. Across the three hospital-based outcomes, the dominant donor is RSV bed-days among adults aged 65 and over in surveillance hospitals — a series that tracks seasonal RSV circulation and severity but is untouched by an infant immunization program. For BM and ICU, adenovirus surveillance and non-RSV respiratory admissions in infants under one year also carry large weights; for medical leave, the leading donors are medical-leave days and bed-days from non-RSV respiratory disease in the same age group. In every case the method selected comparators a clinician would have chosen, rather than arbitrary series that happen to correlate in-sample. Placebo tests with the intervention date reassigned to the 2022 and 2023 seasons recovered no comparable effect.

As shown in Table 3, estimated reductions are: 18,467 days of medical leave ($p = 0.55$), 23,950 outpatient visits ($p = 0.01$), 59,445 basic and intermediate bed-days ($p = 0.01$), and 22,584 ICU bed-days ($p = 0.05$). The medical leave effect is not statistically significant, consistent with the level of year-to-year variation in administrative records (507,339 days in 2023; 483,326 days in 2024); all remaining outcomes exhibit statistically significant reductions.

Of these reductions, the public sector accounts for 20,612 outpatient visits, 51,159 basic/intermediate bed-days, and 19,436 ICU bed-days. These findings underscore the role of the intervention in alleviating pressure on public-sector capacity, which bears the greatest burden of severe respiratory morbidity.

Cohort-level disaggregation reveals that the catch-up cohort, despite slightly lower coverage than the seasonal cohort, accounts for 55.4% of avoided ICU bed-days and approximately half of avoided bed-days (47.9%) and outpatient visits (52.0%). Reductions in medical leave are entirely attributable to the catch-up cohort, as mothers of

seasonal-cohort infants are typically on maternity leave during the winter season. The explanation is one of timing rather than of biology: the catch-up cohort constitutes the whole of the eligible population when the season opens, so its protection is front-loaded onto the epidemic peak, whereas the seasonal cohort accumulates as births occur and contributes progressively through the tail of the season. This makes the catch-up component — absent from most international programs launched in 2024 — the part of the design that relieves the capacity constraint, which is the binding one.

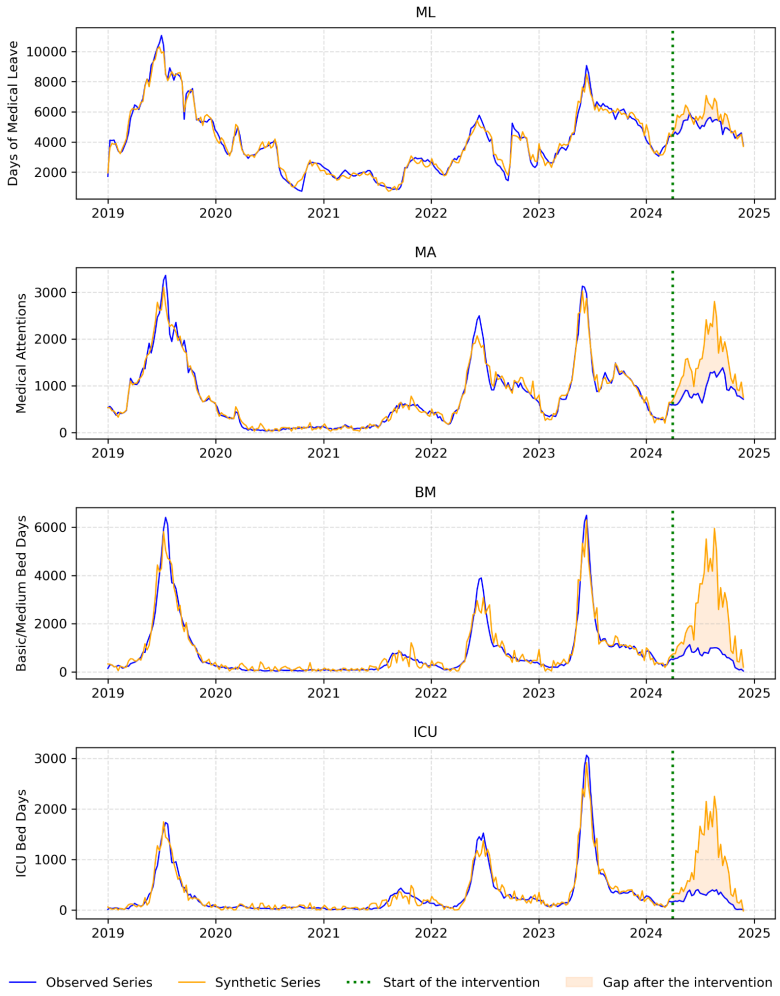


**Fig. 6**: Observed and ASCM-estimated counterfactual outcome series (LRTIs and related viral agents, infants under 1 year of age). Outcomes shown: days of medical leave (ML), outpatient visits (MA), basic/intermediate bed-days (BM), and ICU bed-days. Pre-intervention period: January 2019 – March 2024 (274 weeks). Post-intervention period: April – November 2024 (35 weeks).

**Table 3**: Fitness metrics, $p$-values, and savings estimates for each outcome. ME = Mean Error between synthetic and observed series, reported per week in the pre-treatment row. ML = medical leave days; MA = medically attended outpatient visits; BM = basic/intermediate bed-days; ICU = ICU bed-days. 95% CIs are derived from the distribution of placebo residuals. Dashes mark quantities that are not applicable.

| **Outcome** | **ML** | **MA** | **BM** | **ICU** |
|---|---|---|---|---|
| Pre-treatment ME | −0.53 | −0.61 | −15.61 | −2.88 |
| Post/pre ME difference | −527.1 | −588.3 | −1,682.8 | −642.4 |
| $p$-value | 0.55 | 0.01 | 0.01 | 0.05 |
| Total savings | 18,467 | 23,950 | 59,445 | 22,584 |
| 95% CI | 18,408–37,005 | 21,556–25,106 | 56,323–66,655 | 16,931–23,166 |
| Public savings | 18,467 | 20,612 | 51,159 | 19,436 |
| Catch-up savings % | 100.0% | 52.0% | 47.9% | 55.4% |
| 95% CI | — | 43.1–63.2 | 38.5–60.0 | 41.5–72.0 |
| Seasonal savings % | 0.0% | 48.0% | 52.1% | 44.6% |
| 95% CI | — | 36.8–57.0 | 40.0–61.5 | 28.0–58.5 |

# 5 Impact

## 5.1 Policy impact and program outcomes

The ex-ante analysis delivered to the Ministry of Health in November 2023 provided the economic and epidemiological evidence that justified program approval (Saure et al., 2025). In a decisive budget reallocation, the Ministry secured funding and engaged the manufacturer — which had not planned Chilean market entry before 2028 — to prioritize production for the 2024 season, making Chile the first country to implement a nationwide, universal, free nirsevimab strategy. The program achieved 94% coverage among approximately 155,000 eligible infants across 234 public and private hospitals. Real-world results confirmed 76% effectiveness against hospitalizations, 85% against ICU admissions, and — for the first time in Chile's recorded history — zero RSV infant deaths in both 2024 and 2025.

## 5.2 What the Ministry of Health learned

Three operational insights emerged from the program. First, OR-driven budget reallocation is feasible: the modeling showed that savings from avoided emergency ICU capacity expansion could offset much of the acquisition cost, enabling a budget-neutral reallocation rather than new net expenditure. Second, the real-time monitoring platform established proof-of-concept for integrated hospital surveillance: data from 234 facilities updated daily became a template for future respiratory disease surveillance infrastructure. Third, the shift from a targeted to a universal strategy demonstrated that maximizing equity in access does not compromise cost-efficiency: the universal approach achieved better population-level outcomes at lower per-case cost than the

prior palivizumab-based strategy. The third point is worth stating precisely, because it is counterintuitive. Effectiveness among the high-risk infants that palivizumab had targeted was 79.5%, higher than in the general population, and the universal program reached them at a coverage the targeted program had never achieved. Widening eligibility did not dilute protection where it mattered most; it delivered more of it.

### 5.3 International diffusion and current status

Paraguay modeled its 2025 RSV strategy directly on the Chilean approach, using a parallel cost-saving analysis as scientific justification. Paraguay's 2025 season likewise recorded zero RSV infant deaths. Colombia and Peru are preparing universal strategies for 2026, drawing on the Chilean evidence base. The NirseCL monitoring platform (NirseCL Team, 2024) remains operational, tracking the 2025 RSV season in real time. Two further analyses are currently underway: a longitudinal study of waning immunity across seasons, and a multi-country comparative evaluation of universal versus targeted strategies.

### 5.4 Economic impact

The ASCM counterfactual analysis estimated reductions of 59,000 general hospital bed-days and 22,500 ICU bed-days, generating a net economic benefit of US$20M — including over US$12M in direct savings to the public healthcare system — after accounting for US$45M in acquisition costs. The catch-up cohort alone accounted for 55% of avoided ICU bed-days, reinforcing the importance of including pre-season infants: a design feature absent from most international programs launched in 2024.

## 6 Conclusions

A national immunization program is decided once, in advance, under uncertainty, and then cannot be unwound. What made this instance tractable was not any single technique but the recognition that each stage of the decision required constructing a quantity the data did not contain, and that each construction could be posed as an optimization problem over the same national registries: a code set fitted to sentinel positivity, an efficacy distribution constrained by the ordering that severity implies, a matched sample allocated under scarcity, a season reconstructed from unaffected series. Four methods — integer programming, Monte Carlo simulation, survival analysis, and augmented synthetic control — were deployed not in isolation but as an integrated framework tracking a single policy question from ex-ante design through real-world measurement to ex-post causal evaluation, the pipeline summarized in Figure 1. Three broader lessons emerge for OR practitioners working in public health.

First, timing and credibility matter. Using OR early, before policymakers had committed to any strategy, was decisive. The analytical model, built on national administrative data and validated against sentinel surveillance, convinced both the Ministry of Health and the manufacturer to act. OR's value in policy settings is not only technical: it is reputational and timing-dependent.

Second, validation is the deliverable. In a setting where the estimand is unobservable by construction, the credibility of a number rests entirely on the checks that would have exposed it had it been wrong — a negative control that returns zero, donor weights that a clinician recognizes, placebo tests that find nothing in years when nothing happened, break-even thresholds that state the conditions under which the conclusion fails. Designing the counterfactual analysis as a publicly accessible, prespecified monitoring framework, with data from 234 hospitals updating in real time, made those checks visible to the parties who had to act on them, and made it costly for anyone — including us — to revise the story after the fact.

Third, the framework is scalable and transferable. Paraguay replicated the ex-ante cost-saving analysis for its own context within months, using the same modeling architecture with locally calibrated parameters. The NirseCL framework is not specific to RSV or to Chile. Four inputs are what a country needs to reuse it, one per method: a sentinel surveillance network reporting laboratory-confirmed positivity, against which the code-selection problem of Section 3.1 calibrates a country-specific ICD-10 definition in place of ours; a reported efficacy trial with the same nested severity structure — protection against a milder outcome bounding protection against a more severe one — which the Dirichlet calibration of Section 3.2 requires and which is typical of monoclonal-antibody and vaccine trials; individually linked hospital and immunization registries, without which the matched-sample construction of Section 3.4 has no case–control pairs to allocate; and a donor pool of age–disease series unaffected by the intervention, for the synthetic control of Section 3.5, which a national discharge system with sufficient disease and age granularity will generally provide. Any country with a reliable national health data infrastructure, a seasonal respiratory pathogen, and a new prophylactic agent satisfies enough of these to adapt the framework, even where, as in Paraguay, some inputs must be approximated rather than met exactly. The methods described here, including the Dirichlet-calibrated simulation for uncertain efficacy parameters, the two-stage IP matching for rare subgroup analyses, and the ASCM counterfactual built from administrative donor pools, are directly transferable.

**Limitations.** The framework rests on assumptions that a reader replicating it elsewhere should weigh. Private-sector unit costs were assumed equal to public-sector costs, a conservative assumption only if private costs are in fact higher, as is likely (Section 2). Medically attended visits in the private sector were extrapolated rather than observed, and infant age in that series was imputed from aggregate age bands. The Monte Carlo evaluation of Section 3.2 assumes full and immediate coverage and a constant protective efficacy over a 150-day window; a strategy that fell short of the assumed coverage, or in which efficacy waned within the season, would realize smaller savings than reported. The augmented synthetic control estimates of Section 3.5 identify the effect of the program as a whole and cannot separately attribute it to the individual strategies compared in the ex-ante analysis. Finally, external validity is bounded by what the framework requires: a linked national registry, sentinel virological surveillance, and a donor pool of comparable unaffected series are all necessary inputs, and their absence would require adapting, not simply repeating, the methods described here.

From the Ministry of Health's perspective, the program delivered three strategic improvements over the prior palivizumab-based approach: broader population coverage, elimination of the costly annual temporary ICU capacity expansion required every winter, and a net positive economic return. The absence of RSV infant deaths in both 2024 and 2025 is the most direct measure of impact.

**Acknowledgements.** The Ministry of Health of Chile funded the fieldwork associated with the project. We acknowledge the work of Chile's National Immunization Program, the Departamento de Estadísticas e Información de Salud of the Ministry of Health, the health professionals of all public and private immunization centers in the country, and the national scientific societies that supported dissemination of the strategy. The funders had no role in the design, conduct, analysis, or reporting of this study, nor in the decision to submit it for publication.

## Statements and Declarations

***Funding.***

The Ministry of Health of Chile funded the fieldwork associated with the project. Some of the authors received an independent research grant from Sanofi–AstraZeneca in support of analysis for the studies underlying this work. The funders had no role in the design, conduct, analysis, or reporting of this study.

***Competing interests.***

J. P. Torres reports grants from Sanofi unrelated to the present study. Some of the authors received an independent research grant from Sanofi–AstraZeneca in support of analysis for the studies underlying this work. No other disclosures were reported.

***Relationship to prior publications.***

This manuscript presents the operations research methodology used to design, implement, and monitor Chile's nationwide RSV immunization program. Clinical and epidemiological findings from the program have been reported separately in four medical journals (Saure et al., 2025; Sauré, Torres, et al., 2026; Sauré, Zgheib, et al., 2026; Torres et al., 2025). Those publications focus on clinical and epidemiological outcomes, whereas this manuscript describes the methodology, including the model formulations, calibration procedures, and validation. These methodological details have not been published elsewhere. Any results restated here are included only for context and are appropriately cited. This manuscript has not been published previously and is not under consideration elsewhere.

***Ethics approval and consent to participate.***

The Chilean government routinely collects the data used for this study. Ethical approval was not required because the routine immunization and data collection process was not affected by the realization of this study. Approval for the use of the data, assuring anonymity, was granted by the MoH, whose rigorous ethical standards in data

collection and commitment to maintaining data integrity and anonymity obviated the need for additional ethical approval.

***Data availability.***

The data that support the findings of this study are not publicly available due to confidentiality restrictions. The data were obtained under a data-sharing agreement with the Chilean Ministry of Health that does not permit public dissemination.

***Code availability.***

Code implementing the four analyses described in this paper is available from the corresponding author on reasonable request.

***Author contributions.***

Conceptualization: Leonardo J. Basso, Denis Sauré, Miguel O'Ryan, Juan Pablo Torres; Methodology: Denis Sauré, Marcel Goic, Charles Thraves, Gonzalo Díaz, Natalia Trigo, Amal Zgheib; Software: Felipe del Solar, Ignasi Neira, Amal Zgheib; Validation: Denis Sauré, Marcel Goic, Charles Thraves, Amal Zgheib, Gonzalo Díaz; Formal analysis: Denis Sauré, Marcel Goic, Charles Thraves, Amal Zgheib, Felipe del Solar, Ignasi Neira, Gonzalo Díaz, Natalia Trigo; Investigation: Juan Pablo Torres, Miguel O'Ryan, Felipe del Solar, Ignasi Neira, Amal Zgheib; Resources: Miguel O'Ryan, Juan Pablo Torres, Leonardo J. Basso; Data curation: Felipe del Solar, Ignasi Neira, Amal Zgheib; Writing – original draft: Denis Sauré, Charles Thraves; Writing – review and editing: all authors; Visualization: Felipe del Solar, Ignasi Neira, Amal Zgheib; Supervision: Denis Sauré, Leonardo J. Basso, Miguel O'Ryan, Juan Pablo Torres; Project administration: Leonardo Basso; Funding acquisition: Leonardo J. Basso, Miguel O'Ryan, Juan Pablo Torres. All authors read and approved the final manuscript.

# References


Abadie, A., Diamond, A., Hainmueller, J. (2010). Synthetic control methods for comparative case studies: estimating the effect of california's tobacco control program. *Journal of the American Statistical Association*, *105*(490), 493–505, https://doi.org/10.1198/jasa.2009.ap08746

Ares-Gómez, S., Mallah, N., Santiago-Pérez, M.-I., Pardo-Seco, J., Pérez-Martínez, O., Otero-Barrós, M.-T., ... Martinón-Torres, F. (2024). Effectiveness and impact of universal prophylaxis with nirsevimab in infants against hospitalisation for respiratory syncytial virus in galicia, spain: initial results of a population-based longitudinal study. *The Lancet Infectious Diseases*, *24*, 817–828, https://doi.org/10.1016/S1473-3099(24)00215-9

Ben-Michael, E., Feller, A., Rothstein, J. (2021). The augmented synthetic control method. *Journal of the American Statistical Association*, *116*(536), 1789–1803, https://doi.org/10.1080/01621459.2021.1929245

Bruhn, C.A.W., Hetterich, S., Schuck-Paim, C., Kürüm, E., Taylor, R.J., Lustig, R., ... Weinberger, D.M. (2017). Estimating the population-level impact of vaccines using synthetic controls. *Proceedings of the National Academy of Sciences*, *114*, 1524–1529, https://doi.org/10.1073/pnas.1612833114

Carbajal, R., Boelle, P.-Y., Pham, A., Chazette, Y., Schellenberger, M., Weil, C., ... Enault, M. (2024). Real-world effectiveness of nirsevimab immunisation against bronchiolitis in infants: a case–control study in paris, france, 2024. *The Lancet Child & Adolescent Health*, *8*(10), 730–739, https://doi.org/10.1016/S2352-4642(24)00171-8

Doudchenko, N., & Imbens, G.W. (2016). *Balancing, regression, difference-in-differences and synthetic control methods: a synthesis* (Working Paper No. 22791). Cambridge, MA: National Bureau of Economic Research.

Griffin, M.P., Yuan, Y., Takas, T., Domachowske, J.B., Madhi, S.A., Manzoni, P., ... DeVincenzo, J.P. (2020). Single-dose nirsevimab for prevention of rsv in preterm infants. *The New England Journal of Medicine*, *383*(5), 415–425, https://doi.org/10.1056/NEJMoa1913556

Hammitt, L.L., Dagan, R., Yuan, Y., Baca Cots, M., Bosheva, M., Madhi, S.A., ... Villafana, T. (2022). Nirsevimab for prevention of rsv in healthy late-preterm and term infants. *The New England Journal of Medicine*, *386*(9), 837–846, https://doi.org/10.1056/NEJMoa2110275

Kayvanfar, V., Baldacci, R., Govindan, K. (2025). An integrated approach for enhancing operating room management: capacity planning, fair scheduling, and surgeon resilience. *Annals of Operations Research*, *350*(3), 1385–1412, https://doi.org/10.1007/s10479-025-06565-0

Li, Y., Wang, X., Blau, D.M., Caballero, M.T., Feikin, D.R., Gill, C.J., ... Nair, H. (2022). Global, regional, and national disease burden estimates of acute lower respiratory infections due to syncytial virus in children younger than 5 years in 2019: a systematic analysis. *The Lancet*, *399*(10340), 2047–2064, https://doi.org/10.1016/S0140-6736(22)00478-0

Lopes, J.M., Colon Morales, C., Alvarado, M., Melo, V.A.Z.C., Paiva, L.B., Dias, E.M., Pardalos, P.M. (2022). Optimization methods for large-scale vaccine supply chains: a rapid review. *Annals of Operations Research*, *316*(1), 699–721,

https://doi.org/10.1007/s10479-022-04720-5

Mallor, F., & Azcárate, C. (2014). Combining optimization with simulation to obtain credible models for intensive care units. *Annals of Operations Research*, *221*(1), 255–271, https://doi.org/10.1007/s10479-011-1035-8

Muller, W.J., Madhi, S.A., Seoane Nuñez, B., Baca Cots, M., Bosheva, M., Dagan, R., ... Villafana, T. (2023). Nirsevimab for prevention of rsv in term and late-preterm infants. *The New England Journal of Medicine*, *388*(16), 1533–1534, https://doi.org/10.1056/NEJMc2214773

NirseCL Team (2024). *NirseCL: Real-time rsv monitoring platform.* Retrieved from https://nirse.isci.cl/ (Accessed: 2025-05-14)

Rosenbaum, P.R. (1989). Optimal matching for observational studies. *Journal of the American Statistical Association*, *84*(408), 1024–1032, https://doi.org/10.1080/01621459.1989.10478868

Samii, A.-B., Pibernik, R., Yadav, P., Vereecke, A. (2012). Reservation and allocation policies for influenza vaccines. *European Journal of Operational Research*, *222*(3), 495–507, https://doi.org/10.1016/j.ejor.2012.05.003

Saure, D., O'Ryan, M., Torres, J.P., Trigo, N., Diaz, G., Goic, M., ... others (2025). Cost-savings and health impact of strategies for prevention of respiratory syncytial virus with nirsevimab in chile based on the integrated analysis of 2019–2023 national databases: a retrospective study. *Journal of Infection and Public Health*, *18*(4), 102680, https://doi.org/10.1016/j.jiph.2025.102680

Sauré, D., Torres, J.P., Goic, M., Thraves, C., Pacheco, J., Burgos, J., ... others (2026). Nirsevimab in high-risk infants in a respiratory syncytial virus prevention strategy. *JAMA Network Open*, *9*(4), e266042, https://doi.org/10.1001/jamanetworkopen.2026.6042

Sauré, D., Zgheib, A., Torres, J.P., Goic, M., Thraves, C., Pacheco, J., ... others (2026). Health care utilization and cost implications of chile's 2024 nirsevimab strategy for rsv prevention: a counterfactual analysis. *The Lancet Regional Health–Americas*, *59*, , https://doi.org/10.1016/j.lana.2026.101475

Shi, T., McAllister, D.A., O'Brien, K.L., Simões, E.A.F., Madhi, S.A., Gessner, B.D., ... Nair, H. (2017). Global, regional, and national disease burden estimates of acute lower respiratory infections due to syncytial virus in young children in 2015: a systematic review and modelling study. *The Lancet*, *390*(10098), 946–958, https://doi.org/10.1016/S0140-6736(17)30938-8

Simões, E.A.F., Madhi, S.A., Muller, W.J., Atanasova, V., Bosheva, M., Cabañas, F., ... Villafana, T. (2023). Efficacy of nirsevimab against respiratory syncytial virus lower respiratory tract infections in preterm and term infants, and pharmacokinetic extrapolation to infants with congenital heart disease and chronic lung disease: a pooled analysis of randomised controlled trials. *Lancet Child Adolescent Health*, *7*(3), 180–189, https://doi.org/10.1016/S2352-4642(22)00321-2

Torres, J.P., Sauré, D., Goic, M., Thraves, C., Pacheco, J., Burgos, J., ... others (2025). Effectiveness and impact of nirsevimab in chile during the first season of a national immunisation strategy against rsv (nirse-cl): a retrospective observational study. *The Lancet Infectious Diseases*, *25*(11), 1189–1198, https://doi.org/10.1016/S1473-3099(25)00233-6

Zubizarreta, J.R. (2012). Using mixed integer programming for matching in an observational study of kidney failure after surgery. *Journal of the American Statistical Association*, *107*(500), 1360–1371, https://doi.org/10.1080/01621459.2012.703874

# Appendix A ICD-10 code selection: fit of candidate sets

Figure A1 shows the weekly RSV positivity series implied by each of the three candidate ICD-10 code sets (RSV0, RSV1, RSV2), together with the aggregate positivity rate reported by the sentinel surveillance program (ISP). RSV1 tracks the ISP series most closely across all three years, confirming the solution of the integer program (1) in Section 3.1.

# Appendix B Dirichlet calibration: parameter estimates and fit

Table B1 reports the Dirichlet parameters obtained from the calibration problem (2) of Section 3.2, together with the resulting fit to the MELODY trial efficacy estimates (Muller et al., 2023). The calibrated distribution reproduces the reported means

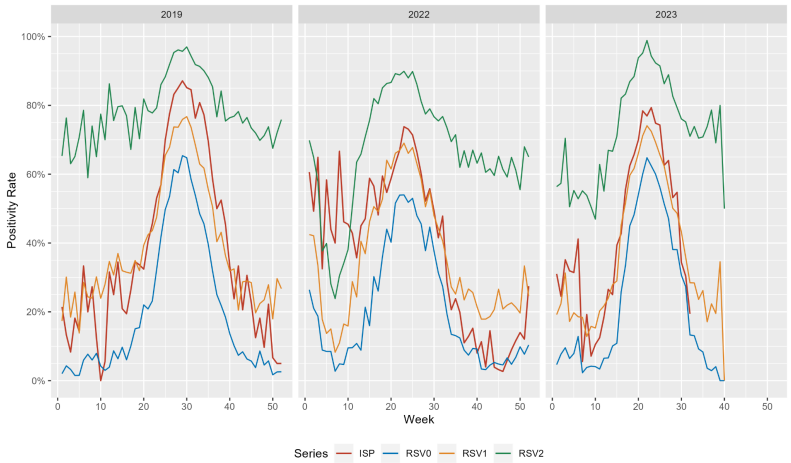


**Fig. A1**: Weekly RSV positivity rates implied by the three candidate ICD-10 code sets (RSV0, RSV1, RSV2) and the aggregate positivity rate reported by the sentinel program (ISP). RSV1 achieves the lowest sum of squared deviations across all three seasons.

exactly and the reported 95% confidence intervals closely, with some narrowing at the lower end for HA and ICU: the nesting constraint that efficacy against a more severe outcome cannot fall below efficacy against a less severe one rules out the lower tail that the independently reported trial intervals admit.

**Table B1**: Calibrated Dirichlet parameters and fit to the MELODY trial efficacy estimates. Components refer to the four-category parameterization of Section 3.2: the first component is efficacy against MA, the second and third the increments from MA to HA and from HA to ICU, and the fourth the residual. MA = medically attended; HA = hospital admission; ICU = intensive care unit.

| *Panel A: calibrated parameters* | | | | |
|---|---|---|---|---|
| | **MA** | **HA − MA** | **ICU − HA** | **Residual** |
| $\alpha_j$ | 39.60 | 0.21 | 0.93 | 11.09 |
| *Panel B: efficacy, % (95% CI)* | | | | |
| | **MA** | **HA** | **ICU** | |
| Reported (MELODY) | 76.4 (62.3–85.2) | 76.8 (49.4–89.4) | 78.6 (48.8–91.0) | |
| Calibrated | 76.4 (64.1–86.8) | 76.8 (64.5–87.1) | 78.6 (66.6–88.5) | |

# Appendix C Augmented synthetic control

Consider $N$ weekly time series, where series 1 is the outcome of interest (e.g., infant ICU bed-days attributable to LRTI) and series $2, \ldots, N$ constitute the donor pool of 266 age–disease groups. Let $y_{it}$ denote the value of series $i$ in week $t$. The intervention begins at week $T_0$, and the post-intervention window spans weeks $T_0 + 1, \ldots, T$.

Because the donor pool contains 266 series — nearly as many as the $T_0 = 274$ pre-intervention weeks — a direct application of the standard synthetic control method (SCM) would overfit. We augment the model with $K$ covariates and $M$ transformations of the outcome (e.g., cumulative monthly cases, seasonal indicators), yielding extended observation vectors

$$z_{it} = \big(y_{it},\ x_{it1}, \ldots, x_{itK},\ y_{it}^{(1)}, \ldots, y_{it}^{(M)}\big) \ \in \ \mathbb{R}^{1+K+M}.$$

Transformation weights $v_j$ are derived from XGBoost variable-importance coefficients when predicting $y_{1t}$ from all series features in the pre-intervention period.

The method proceeds in two steps. First, the SCM solves the constrained regression

$$\mathbf{w}^{\mathrm{SCM}} = \arg \min_{\mathbf{w} \geq 0,\, \mathbf{1}^\top \mathbf{w} = 1} \ \sum_{j=1}^{1+K+M} v_j \sum_{t=1}^{T_0} \big(z_{1tj} - \textstyle\sum_{i=2}^{N} z_{itj}\, w_i\big)^2. \tag{C1}$$

If the pre-treatment fit is good, $\mathbf{w}^{\mathrm{SCM}}$ provides unbiased counterfactual weights. With a large donor pool, however, the convex-hull restriction $w_i \geq 0$ may preclude a good fit. ASCM (Ben-Michael et al., 2021) therefore relaxes non-negativity and applies a ridge-regression bias correction,

$$\mathbf{w}^{\mathrm{AUG}} = \arg \min_{\mathbf{w}:\, \mathbf{1}^\top \mathbf{w} = 1} \left\{ \frac{1}{2\lambda} \sum_{t=1}^{T_0} \Big(y_{1t} - \textstyle\sum_{i=2}^{N} y_{it}\, w_i\Big)^2 + \frac{1}{2} \sum_{i=2}^{N} \big(w_i - w_i^{\mathrm{SCM}}\big)^2 \right\}, \tag{C2}$$

which is equation (6) of Section 3.5 stated in full. Here $\lambda > 0$ controls the trade-off between pre-treatment fit (first term) and proximity to the SCM baseline (second term). The second term covers only the original outcome series $y_{1t}$, not the augmented features, to avoid overfitting the transformations. The parameter $\lambda$ was selected by cross-validation over a grid, using placebo tests in the pre-intervention period, yielding $\lambda = 10$ for all four outcomes (ML, MA, BM, ICU).

Statistical significance is assessed via placebo tests: the same estimation procedure is applied to each donor series in turn, generating a distribution of placebo treatment effects. The $p$-value for the true treatment effect is the proportion of placebo effects, in absolute value, exceeding the observed effect (Bruhn et al., 2017). Confidence intervals for cumulative savings are obtained from the distribution of cumulative placebo residuals in the post-intervention period. The cohort-level decomposition (catch-up vs. seasonal) is performed with additional Monte Carlo simulations that allocate counterfactual cases proportionally across cohorts.