Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

TARGETED HOUSEHOLD QUARANTINING: ENHANCING THE EFFICIENCY OF EPIDEMIC RESPONSES

Johannes Ponge[1], Julian Patzner[2], Bernd Hellingrath[1], and André Karch[3]

[1]Dept. of Information Systems, University of Münster, Münster, NRW, GERMANY
[2] IMEBI, Martin Luther University Halle-Wittenberg, Halle (Saale), ST, GERMANY
[3]Inst. of Epidemiology & Social Medicine, University of Münster, Münster, NRW, GERMANY

ABSTRACT

Non-pharmaceutical interventions (NPIs) are the immediate public health reaction to emerging epidemics. While they generally help slow down infection dynamics, they can be associated with relevant socioeconomic costs, like lost school- or work days caused by preemptive household quarantines. However, research suggests that not all households contribute equally to the overall infection dynamics. In this study, we introduce the novel "Infection Contribution" metric that allows us to trace the involvement of particular household types over entire infection chains. Building upon the German Epidemic Microsimulation System, we quantify the impact of various household types, considering their size and composition in a COVID-19-like scenario. Additionally, we show how targeting interventions based on household characteristics produces efficient strategies, outperforming non-selective strategies in almost all scenarios. Our approach can be transferred to other NPIs, such as school closure, testing, or contact tracing, and even inform the prioritization of vaccinations.

1 INTRODUCTION & BACKGROUND

The spread of infectious diseases remains a major global health challenge. Recent pandemics such as COVID-19 have highlighted the need for effective containment measures. Non-pharmaceutical interventions (NPIs), such as social distancing, mask mandates, or isolation rules, are essential when pharmaceutical solutions are unavailable. While these interventions can slow transmission and reduce the number of infections and deaths (Aleta et al. 2020; Geffen and Low 2020; Weigl et al. 2021; Zhang et al. 2022), they often have significant social and economic costs, such as disruptions to education, reduced workforce availability, mental health impacts, and income loss (Chatterjee and Chauhan 2020; Chen et al. 2024; Giallonardo et al. 2020; Jin et al. 2021).

As prominently applied during the COVID-19 pandemic, universal quarantining of households with symptomatic individuals or potentially infectious contacts is often one of the immediate responses to an outbreak scenario. However, research has found that the impact of households on the force of infection depends heavily on their internal structure and the distribution of particular household characteristics within a population. The number of large households, for example, is an important factor in the overall disease dynamics (Dönges et al. 2024; Liu et al. 2021), given that household size significantly informs the likelihood of within-household infections that subsequently propagate the disease into the community (Møgelmose et al. 2023). The distribution of household sizes in a region can, therefore, have a significant influence on the spread dynamics (Liu et al. 2021). Moreover, the number of schoolkids and working adults in a household is an important criterion, as both groups may have an increased likelihood of bringing infections into the household (Endo et al. 2019; House et al. 2022; Møgelmose et al. 2023; Tseng et al. 2023). Schoolkids often have high contact rates at schools (Grijalva et al. 2015; Guo et al. 2023; Tseng et al. 2023), and schools themselves can be considered contact hubs for large numbers of families. This makes households with many schoolchildren potential hotspots for infections. Similarly, working adults may be

at increased risk of becoming infected and carrying the virus into their household (Nash et al. 2022; Potter et al. 2015). Considering these heterogeneities, two questions arise: (1) What is the impact of certain household types on the disease dynamics on a macroscopic scale? And furthermore, (2) Can intervention strategies target high-impact households to better balance disease mitigation with social costs?

Several related modeling studies have been published in recent years. Dönges et al. (2024) developed an ODE-SIR model to analyze disease transmission within and between households of different sizes. They show the importance of incorporating data about the household size distribution in a population for accurate calculations of the basic reproduction number. They do, however, not consider the individual composition of households. Endo et al. (2019) provide useful insights into the dynamics of influenza at the household level, emphasizing that understanding household-specific contact patterns helps to illustrate how prevention measures should be adapted. Møgelmose et al. (2023) used longitudinal microdata from Belgium to explore how developing demographic structures condition the spread of diseases using an agent-based model. The authors show how age-dependent household compositions significantly impact overall disease dynamics. While the authors suggest that these findings can support the development of targeted intervention strategies, they do not explicitly simulate intervention scenarios.

Our paper aims to quantify the impact of various household sizes and compositions on the overall infection dynamics and explore whether interventions targeting specific household types can result in more efficient strategies concerning the slowdown of infections and incurred quarantine days. To achieve these goals, we introduce a novel Infection Contribution metric to quantify how different household types, based on household characteristics, contribute to the propagation of infections within heterogeneous populations. We apply this metric to a COVID-19-like disease outbreak scenario that we simulate using an agent-based model of the Saarland region in Germany. Subsequently, we repeat the outbreak scenario but apply targeted intervention strategies to households based on size and composition. We identified a selection of efficient quarantine strategies that balance epidemic control with societal cost with respect to lost school- and workdays. While we conducted this analysis for household quarantine measures, we suggest that the framework is adaptable to other NPIs, such as school closure, testing, or contact tracing, and could even inform the prioritization of vaccinations.

Chapter 2 introduces the general population and disease model. Chapter 3 presents the novel metric and discusses the importance of different household types on the overall disease dynamics. We simulate the targeted quarantine strategies in Chapter 4 and offer a brief discussion and conclusion of our work in Chapter 5.

2 POPULATION & DISEASE MODEL

The simulations in this study are carried out using the German Epidemic Microsimulation System (GEMS), which was first introduced by Ponge et al. (2023). GEMS is a discrete-time agent-based infectious disease modeling framework with a full-scale virtual population of all German states. These population models are based on the Gesyland project. They are generated using census information and contain realistic demographic structures, both on an individual level (age, sex, employment status, etc.) and household level (household size, family type, etc.). We use the Saarland model for our experiments as it is the smallest German area-state (as opposed to city-states with vastly different demographic structures) with roughly one million individuals. The agents in GEMS are associated with so-called "settings". They represent physical or social contexts in which person-to-person contacts are realized, potentially leading to an infection. From a technical perspective, they can be regarded as a simple subset of agents. All agents have a household setting and, depending on their age and employment status, an associated workplace or school setting, each with an internal hierarchical structure. Workplaces are split into departments and offices; schools have school years and classes. A peculiarity of the Gesyland population is that the school setting also covers nursery schools and universities. The household setting also covers care homes, leading to a small number of very large households comprising almost entirely elderly people. Moreover, all agents are associated with a municipality setting based on the geolocation of their household. These settings simulate semi-random non-household-, non-school-, non-workplace-encounters in geographical proximity. Figure 1

illustrates the geolocations and the size distributions of households, workplaces, and schools. Moreover, the map shows municipality borders.

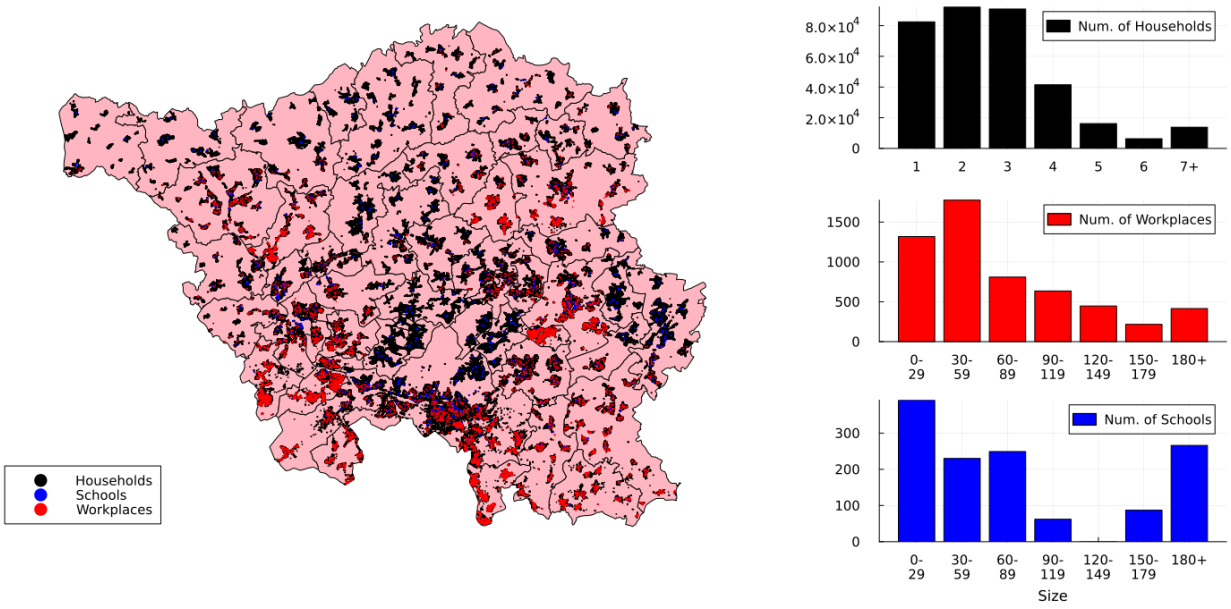

Figure 1: Population model and setting size distribution.

We assume an average daily contact rate for each setting type following the POLYMOD contact survey (Mossong et al. 2008). There are 3.3 daily contacts in households, 2.8 contacts in workplaces (split into workplace, department, and office level), 1.8 school contacts (split into school, school year, and class level), and 4.0 random outdoor contacts in the surrounding municipality. We parameterize a SARS-CoV-2-like pathogen (wild type) that follows a basic SEIR model. The average onset of symptoms is 6.55 days after exposure (Wu et al. 2022), although agents become infectious 2 days earlier (Byrne et al. 2020). The average time to recovery after symptom onset is 13.4 days (Byrne et al. 2020). Contact rates and disease timings are implemented as expected values of a Poisson distribution. Based on Byambasuren et al. (2020), we simulate 17% of agents to have an asymptomatic disease progression. Moreover, we assume that 0.2% of agents aged 0-39 years, 0.9% of agents aged 40-59 years, 5% of agents aged 60-79 years, and 14.8% of agents aged over 79 years will die from the disease (Verity et al. 2020). We use the per-contact transmission rate β to fit the basic reproduction number 3.28 (Liu et al. 2020). All scenarios start with 0.1% of randomly infected individuals and simulate one year. Figure 2 shows the overall disease progression in the unmitigated baseline scenario.

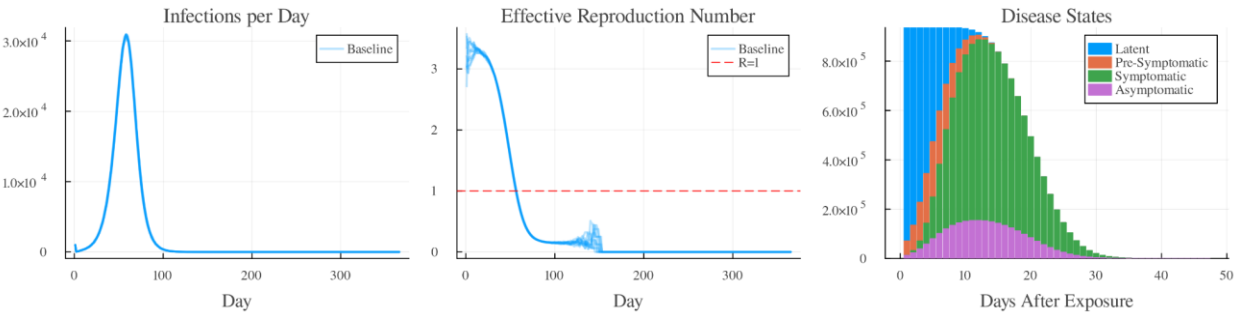

Figure 2: Ten runs of the unmitigated disease progression starting with 0.1% initial infections.

With an initial reproduction number of 3.28 and an average generation time of roughly 9 days, the unmitigated epidemic runs through the population within the first three months. Eventually, 96.3% of the population will be infected. On average, the infection lasts for 20 days from exposure to full recovery.

3 CONTRIBUTION OF HOUSEHOLD TYPES TO OVERALL INFECTION DYNAMICS

3.1 Definition of Household Types

As mentioned in Chapter 1, the literature suggests that households of varying sizes and compositions contribute differently to the overall infection dynamics (Dönges et al. 2024; Endo et al. 2019; Liu et al. 2021). Therefore, the household types we focus on in our analysis each consist of a size attribute (any, 2+, 3+,...,6+) and a composition attribute (any, with school kids, with 2+ school kids, without schoolkids, kids in multiple schools, 1+ kid in a 150+ member school, with workers, with 2+ workers, without workers, with workers and school kids, without schoolkids but with workers, and without school kids or workers). This leads to a total of 72 household types (combinations). Size attributes state the minimal size, not the exact size, as we intend to derive actionable strategies, and quarantining two-person households only seems impractical. The composition attributes target the presence of school kids or working adults in the respective households. Most household types are not mutually exclusive.

3.2 Infection Contribution Metric

We aim to understand the relative importance of different household types to the infection dynamics. To achieve that, we introduce a new metric called *Infection Contribution (IC)* that provides information about the importance of any given infection with respect to the overall infection dynamics. We then aggregate the results for all infections caused by people living in a particular household type.

The Infection Contribution quantifies the number of direct and indirect descendant cases for each infection. We calculate the number of secondary infections to which any index infection *contributed*. This metric has several advantages over person- or setting-dependent R-values, as it works independently from the underlying network structure. Let us consider an extreme example of two disjunctive subnetworks and an infectious pathogen introduced to one subnetwork but not the other. Moreover, we assume a single edge connecting one individual from each subgraph. If that connection leads to the introduction of the pathogen to the fully susceptible subnetwork, this infection will have a vast IC value, even if the person's reproduction rate is just 1 (assuming the infector did not infect anybody else). The IC metric can consider what happens "down the road" and identify high-impact traits in the model.

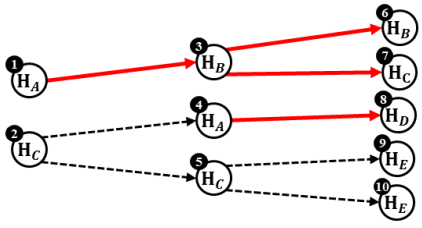

Infection	Contribution
1	3
2	5
3	2
4	1
5	2
6 – 10	0

Household Type	Contribution	Contribution Ratio
H_A	4	50%
H_B	2	25%
H_C	5	62,5%
H_D	0	0%
H_E	0	0%

Figure 3: Calculation of infection- and household-type contributions.

Figure 3 shows an example of an infection tree with two seeding infections. The top-left number in black circles indicates the infection ID; the white circles indicate the infected person's household type (A-E). The center table shows the cumulative number of subsequent cases that can be traced back to the respective infection. As seen in the graph, infections 3, 6, and 7 are successors of infection 1, resulting in a contribution of 3.

To understand the relative contribution of a particular household type, we calculate for each household type H its contribution ratio C_H as the cardinality of the union of infection sets S_i originating from infections $i \in I_H$ of people living in household type H, normalized by the total number of infections N:

$$C_H = \frac{\left|\bigcup_{i \in I_H} S_i\right|}{N},$$

where I_H is the set of infections of people associated with household type H. The union operator is important to prevent double counts in infection chains that contain the same household type multiple times. In Figure 3, infections 1 and 4 are caused by individuals who live in household type A, resulting in a contribution value of 4 for this type. With eight secondary infections (excluding the seeds), we get a household contribution rate of 50% for type A. In other words, 50% percent of all infections can be traced back to a person living in a type A household. A variant of this metric can be set up to identify the contribution of infections and household types to cases that resulted in death. Both metrics identify households that disproportionately drive transmission or severe outcomes, providing insights for targeted interventions.

3.3 Calculating Infection Contributions in the Unmitigated Scenario

For this analysis, we ran the contribution metric from Chapter 3.2 on each household type identified in Chapter 3.1 for the unmitigated baseline scenario for the Saarland model to understand which household types are the strongest drivers of the overall infection dynamics. Figure 4 illustrates the contribution ratio, the percentage of infections involving a particular household type at any stage in their infection chain if traced back to one of the initial cases (Y-axis). Moreover, as the goal is to derive efficient intervention strategies, the X-axis shows the fraction of the population that lives in a household of that type. The subplots are stratified by the minimum household size. The shapes represent different compositions (school-focused compositions in red, work-focused compositions in blue, and combinations in orange). Generally speaking, a dot in the bottom-left suggests *few people are causing few infections*, and *many people are causing many infections* in the top-right. To develop targeted intervention strategies, the sweet spot is the top left, where *few people cause many infections*.

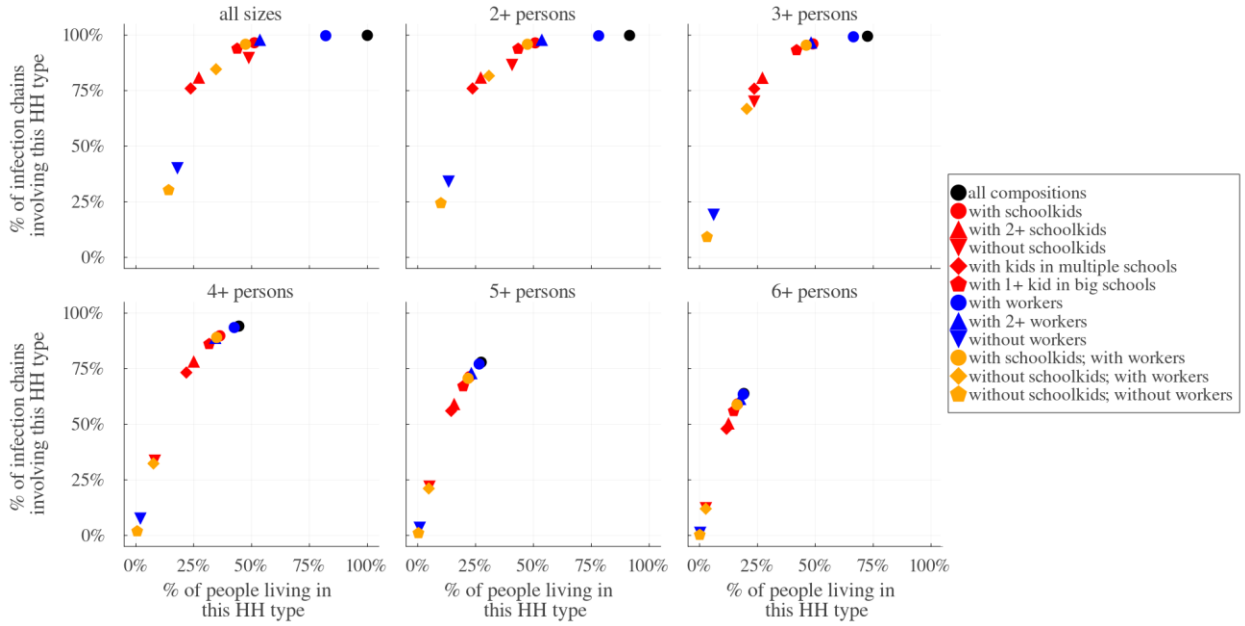

Figure 4: Contribution and affected individuals per household type.

All household types have a disproportionate contribution compared to the fraction of people living in the respective household type. We expect this as the chance of a household type being included in any given infection chain increases exponentially with the length of the chain. However, larger households in our model generally have a higher contribution-to-affected-people (CP) ratio, regardless of their composition. While households of a minimum size of two are part of 99.9% of all infection chains, only 91.5% of people live in such a household, resulting in a CP-ratio of 1.09. Three-or-more-person households have a 1.37, four-or-more-person households a 2.12, five-or-more-person households a 2.84, and six-or-more-person households a 3.34 CP-ratio. In fact, if we order all household types by their CP-ratio, the first 21 entries belong to households of size four and above. While this might hint towards the efficiency of intervention measures regarding induced quarantine person days, one must also consider that larger households are part of fewer infection chains overall (e.g., 63.8% for six-or-more-person households). This poses a hard limit, even for a "perfect" intervention strategy.

Another interesting observation is that the contribution rates of households with schoolkids (red circle) and households without schoolkids (red downward triangle) are very similar when all household sizes are considered. A similar contribution value is understandable, given that each household type covers roughly 50% of the population. Based on the body of literature on school kids being drivers of infections (Grijalva et al. 2015; Guo et al. 2023; Tseng et al. 2023), however, we would have expected a more significant difference. With bigger minimum household sizes, the gap in contribution widens, coinciding with the fact that there are fewer large households without schoolkids. The CP-ratio, e.g., for three-or-more-person households, diverges noticeably (1.96 with and 2.97 without schoolkids). The situation is very different when looking at households with and without workers (blue circle and downward triangle). We see a significant gap for all household sizes, both in contribution and affected people. The graphs show two clusters where, for household sizes 1-3 and above, the two household types that explicitly exclude workers are far below the 50% mark concerning their contribution. For sizes 4-6 and above, the types explicitly excluding schoolkids move from the upper to the lower cluster. Coincidentally, larger households (4-6+) without schoolkids have the highest CP-ratio (above 4.2) while being among the rarest household types in the model.

Focusing only on households with two or more schoolkids and/or workers (upward triangles) makes a considerable difference, especially for the smaller minimum household sizes. While only half of the people who live in a household with schoolkids live in one with two or more children, they still have a comparatively high contribution (80.1% compared to 96.6%). This increases those households' CP-ratio (across all sizes) from 1.89 to 2.97. The effect is even more pronounced when comparing households with workers and those with two or more workers, where the overall contribution is very similar (99.6% and 97.9%), but the number of people living in such households differs by 28.4%.

Whether or not one schoolkid is assigned to a 150+ sized school causes a 13% increase in the CP-ratio over households with *any* schoolkids assigned. Households with kids in multiple different schools scored a 7% increase over households with at least two schoolkids in *any* school. Both effects appear to be moderate.

4 SIMULATION OF TARGETED INTERVENTION STRATEGIES

4.1 Household Quarantine Strategy

The previous chapter highlighted how different household types have significant variations in their contribution to the overall infection dynamics. This inevitably leads to the question of whether that knowledge can steer the development of targeted household quarantining strategies.

There are, of course, numerous ways to design quarantine strategies in practice. To reduce complexity and generate comparable results, we focused on one particular setup where a household must undergo a 14-day quarantine once any member experiences symptoms. All individuals are perfectly compliant.

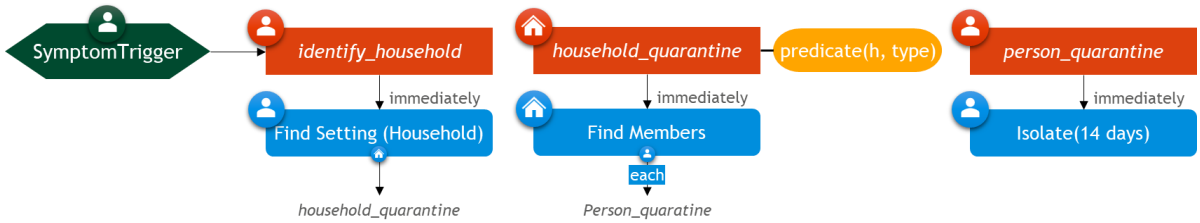

Figure 5: Household quarantine intervention scenario in TriSM formalization.

The strategy is visualized in Figure 5, following the Trigger-Strategy-Measure (TriSM) formalization (Ponge et al. 2024). Once an individual experiences symptoms, a household identification measure is executed, discovering the associated household setting and putting every member into quarantine for two weeks. This will effectively prevent all out-of-household contacts for that duration. Any further household member developing symptoms resets the two-week period for all members. The orange "condition"-box indicates that the quarantine will only apply if the household matches certain criteria, e.g., being of a particular type (like a three-person household with schoolkids). For each of the 72 household types introduced in Chapter 3.1, we ran ten simulations where the quarantining strategy only applied to households of the respective type.

4.2 Simulation Results

We ran the simulations on a server with two 12-core Intel Xeon Gold 6136 CPUs and 128GB of system memory. Each run, simulating one year with one million agents, took roughly 20 seconds. While we focused on a particular German state, runtime and memory usage in GEMS scale close to linearly with the population size. The full German population model can be run on the same machine in about 30 minutes.

Figure 6 contrasts the strategies' benefits and costs. We measure the strategy benefit as the reduction in the basic reproduction number R0 compared to the unmitigated baseline scenario presented in Chapter 2. Since we focus on isolation strategies, the costs are defined as the total person-days spent in quarantine. Moreover, we stratify the results by lost school days (caused by quarantined school kids) and lost work days (caused by quarantined workers). Colored dots indicate that the strategy is Pareto-optimal, meaning no other strategies have higher benefits at lower costs. The colored dots make up the efficiency frontier. The dot sizes represent the strategies' household size limit (largest dots: any household size, smallest dots: 6+ person households only). Grey dots are strategies that are not Pareto-optimal and are all equally sized.

Due to population size and the almost 1,000 seeding infections, the results are relatively stable across simulation runs. Initial experiments demonstrated that ten repetitions per scenario offer an effective balance between result accuracy and computational efficiency. The maximum coefficient of variance (CV) in R0 across all scenarios is 0.05, meaning that the highest standard deviation among the ten repetitions in any of the scenarios is 5% of their mean value for R0. The maximum CV for total infections is 0.2%, 0.9% for total deaths, 2.7% for total quarantine days, 3.4% for lost work days, and 0.6% for lost school days.

As expected, we observe a strong general correlation between the number of person-days spent in isolation and the reduction in R0 in the top left graph of Figure 6. However, we also notice that all scenarios are not equally efficient, meaning that there are strategies that yield higher benefits (R0 reduction) at lower costs (quarantine days). In fact, quarantining all households, regardless of their composition (black dots), is not an optimal solution in most cases. It is only an efficient option if it applies to all household sizes, as this is the most effective solution overall. Considering the household composition for targeted measures seems reasonable in all other situations. We also notice that the slope of the efficiency frontier increases the closer scenarios get to the phase transition (R0 = 1, i.e., a reduction of 2.26 given our initial parameterization), where the epidemic would become subcritical and die out. We ran experiments with lower initial R0-values where some of the quarantining scenarios were able to push R below 1, and in those cases, the results are clear: the strategy that ends the epidemic the soonest is both the most efficient and most effective. If the epidemic cannot be *ended,* it has to be *managed*.

An interesting observation is that eleven out of the 25 Pareto-optimal scenarios focus on households of various sizes that explicitly exclude school children (orange diamonds and pentagons, red downward triangles). Among those, the maximum possible R0 reduction is 0.74 when quarantining households of all sizes without school children. Above that, the best results are yielded by households (sizes 2+ and 3+) with one or more children in a big school (red pentagon), households (sizes 1+, 2+, and 3+) with schoolkids and workers (orange circle), and, especially, households (sizes 1+ and 2+) with two or more workers (blue upward triangle). The latter marks the most effective (R0-reduction of 1.12) option that does not require quarantining everybody. The other end of the spectrum shows that eight efficient solutions focus on larger households (4+, 5+, and 6+) of various compositions. However, the largest reduction in R0 we can achieve focusing on these households is 0.25 when targeting 5+ sized households with kids in multiple schools (red diamond). Concerning the correlation between household size and effectiveness, the blue downward triangle for households of any size without workers makes an exception. Targeting these households will only yield an R0 reduction of 0.19.

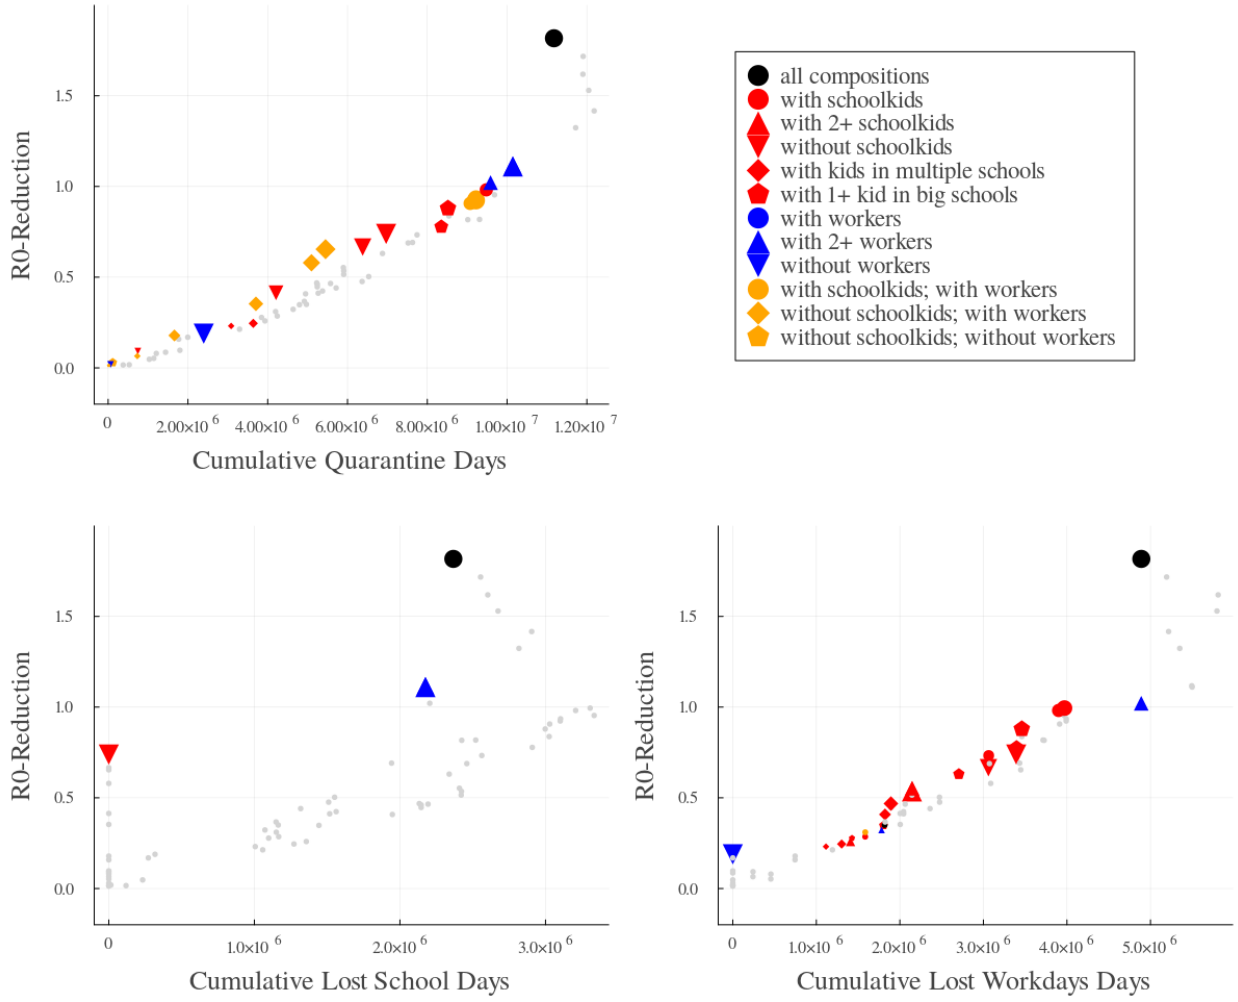

Figure 6: Efficiency frontier of targeted household quarantine strategies regarding their reduction in the basic reproduction number R0 (compared to the unmitigated scenario in Chapter 2) and induced total quarantine days (top left), lost school days (bottom left), and lost work days (bottom right). Colored dots represent Pareto-optimal strategies.

Only four scenarios are Pareto-optimal when solely considering lost school days as costs (Figure 6, bottom left). Self-evidently, any scenario that explicitly excludes schoolkids has zero costs. Among those, the scenario focusing on households of any size without children (red downwards triangle) has the highest benefit (R0 reduction of 0.74). Surprisingly, this is also an efficient scenario concerning lost work days (Figure 6, bottom right). In fact, it is the only scenario (other than the all-sizes-all-compositions scenario) that is Pareto-optimal in all three analyses. The two remaining options are scenarios that focus on 1+ and 2+ sized households with two or more workers (blue downward triangles on top of each other).

The overall graph looks much different when only considering lost work days (Figure 6, bottom right). There are also a few zero-cost options, with the scenario focusing on households without workers of any size (blue downward triangle) being the most effective. However, the achievable reduction in R0 is narrow, given that the number of households without workers in the baseline population is small. Most efficient scenarios (17 out of 24) target households containing schoolkids in various compositions (red dots in diverse shapes). Optimizing for lost work days seems to "push" the costs of quarantining towards children. There are two noteworthy exceptions, as strategies targeting households of sizes 1+ or 2+ without schoolkids are also efficient options with respect to lost work days.

4.3 Comparing Quarantine-Day Utilization and CP-Ratio

In Chapter 3.3, we evaluated which household types have a disproportionately large contribution to the overall infection dynamics compared to the number of people living in a household of that type (CP-ratio). We now contrast these values with the number of person-quarantine days that were incurred to reduce the basic reproduction number by 0.1 per scenario, which we derived from the simulations in Chapter 4.2. We initially assumed that higher CP-ratios must yield a lower quarantine-day-per-R0-reduction value, as we would target "the right" individuals. The results, however, are counterintuitive.

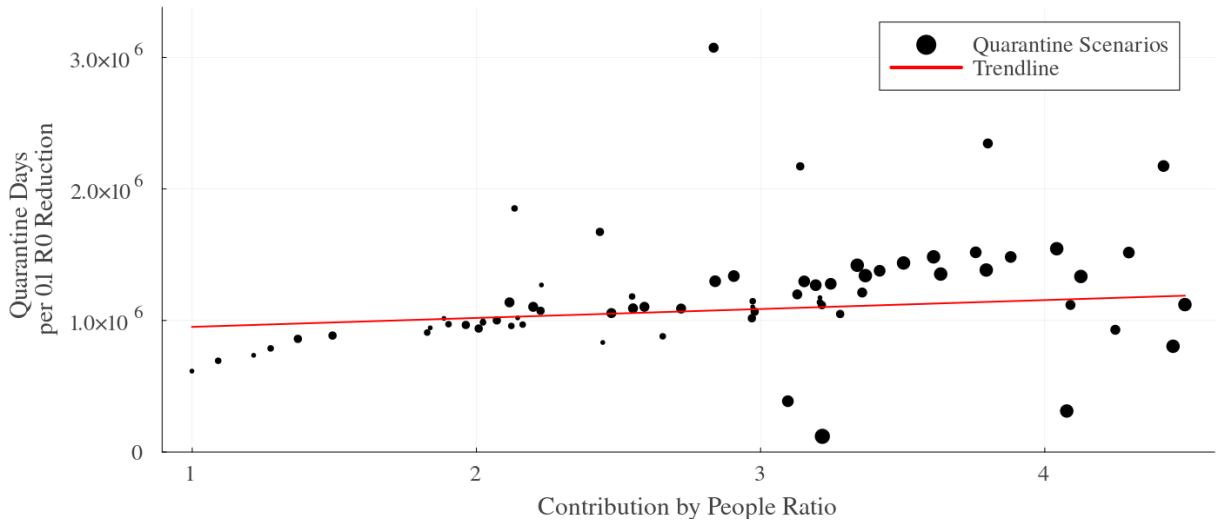

Figure 7: Quarantine-day efficiency and contribution by affected people ratio. Dot sizes represent the minimum size of the respective household size (1+ to 6+).

Figure 7 shows a slightly positive correlation between the CP-ratio and the number of quarantine days, which means that the higher the CP-ratio, the less efficient each quarantine day will be. We offer two potential explanations. First, household types with the highest CP-ratios have the lowest absolute contribution. Even without these households, the overall force of infection is still high. Removing a certain household type from the infection chains only prevents individuals from being infected by someone living in *that* household type. Still, it does not prevent them from getting infected *at all*, given that the underlying population model is generally well-connected. The second reason is rooted in the intervention strategy

design. It mandates symptomatic individuals to stay home, preventing out-of-household infections but not within-household infections. As higher CP-ratios correlate with larger households (see dot sizes in Figure 7), we suggest quarantine days become less effective since they do not prevent in-household infections. Our observations coincide with the suggestion of Dönges et al. (2024) that household quarantine measures are more effective in populations with smaller average household sizes. Future research should analyze whether the results would look similar if individuals were quarantined in a dedicated facility.

5 DISCUSSION & CONCLUSION

The analyses show that calculating the Infection Contribution of individuals or households with particular characteristics can be a valuable approach to targeting interventions. While the calculation of R-values is useful for estimating the pace of infection dynamics, outbreak size, and herd immunity thresholds, the IC metric can identify drivers of infection dynamics. It is not intended to replace R calculations but to provide an addition to the epidemiological toolset. As it requires complete knowledge about the infection tree, the IC metric is mainly a theoretical construct, suitable for (individual-based) simulations. In our experiments, we focused on a baseline scenario with a large population, a considerable number of initial infections (0.1%), and a moderately infectious pathogen. Under these conditions, the IC metric provides stable results across all runs. Nonetheless, like all agent-based models, the IC metric is susceptible to path dependency and therefore fluctuating outcomes for scenarios with very few seeding infections, small populations, or R0 values close to 1. However, in practice, these conditions are not representative of a scenario that requires large-scale targeted intervention planning in any substantial way. The IC metric also works for scenarios with much more infectious pathogens, although the relevance of targeted intervention strategies under those conditions is debatable, as it likely necessitates the most effective, rather than the most efficient strategy.

Our simulations of the federal state of Saarland confirm the conclusions of previous modeling studies (Dönges et al. 2024; Endo et al. 2019; Møgelmose et al. 2023) that households contribute differently to the overall infection dynamics. The largest households in our study consistently show the highest contribution-to-affected-people ratio, meaning that those households are disproportionately often involved in infection chains. Moreover, the results indicate that the composition of households, especially with respect to the presence of schoolkids and working adults, plays a significant role. We can, however, not conclude that targeting household types with high CP-ratios will always yield the most efficient quarantining strategies due to various factors, such as their absolute importance for the infection dynamics and the fact that quarantining does not prevent in-household transmissions. That said, we suspect that these conclusions do not generally apply to all types of interventions. The approaches presented in this work could have an even larger impact when developing targeted vaccination-, testing-, or contact-tracing strategies, especially when considering limited resources (vaccine doses, testing kits, public health personnel) as the cost function instead of quarantine days. With regard to the household quarantine strategies presented here, we notice that strategies targeted towards particular household compositions are generally superior to one-fits-all interventions in almost all scenarios.

Of course, our observations must be taken with a grain of salt, as we used a population model of a particular German federal state. We repeated the complete analysis for the eastern German state of Saxony-Anhalt, which has a similar demographic structure and is also primarily rural. The results were comparable to the Saarland analysis. When we re-run the analysis for the city-state of Berlin, the general trends about the contribution of differently sized households were similar. However, the efficiency frontier of household composition-based quarantining strategies deviated significantly. In subsequent studies, we plan to apply our methodology to the full German population model that is built into GEMS. Moreover, we plan to analyze how the introduction of sub-perfect mandate adherence affects the efficiency frontier. As mentioned before, we also saw substantial differences when we changed fundamental assumptions about the pathogen, such as the basic reproduction number. We thus argue that the evaluation of targeted intervention strategies must be conducted on a per-population and per-pathogen basis.

Whether quarantining households based on specific criteria is applicable in public health practice is debatable and particularly dependent on the regulatory body of the respective region. Moreover, household

quarantining would be only one aspect of a strategy mix (potentially including school closure, travel suspension, contact tracing, and others) applied simultaneously in a real-world situation, depending on how sharply the infection dynamics have to be reduced to not overburden the healthcare system. Future research should also consider strategy combinations. The approach in this study produces a spectrum of efficient options for one NPI class. It enables decision-makers to understand the implications of both the infection dynamics and associated socio-economic costs. Our analysis is primarily an exemplary case study on systematically deriving efficient, targeted intervention strategies. The approach is transferable to a variety of interventions and can not only be used to target high-impact demographics but also to prioritize individuals when working with limited available resources (e.g., vaccine doses or testing kits). Given that producing these simulation results takes around half a day on the hardware we had available, we argue that this approach can support short-term decision-making during emerging epidemics.

The code to replicate this study is available on GitHub.

ACKNOWLEDGMENTS

This work is part of the "OptimAgent" project. The project "OptimAgent" was funded by grants (no. 031L0299X) of the German Federal Ministry of Education and Research (BMBF). "OptimAgent" is part of the Modeling Network for Severe Infectious Diseases (MONID). The authors are responsible for the content of this publication.

REFERENCES

Aleta, A., D. Martín-Corral, A. Pastore Y Piontti, M. Ajelli, M. Litvinova, M. Chinazzi, *et al*. 2020. "Modelling the Impact of Testing, Contact Tracing and Household Quarantine on Second Waves of COVID-19". *Nature Human Behaviour* 4(9):964–971.

Byambasuren, O., M. Cardona, K. Bell, J. Clark, M.-L. McLaws, and P. Glasziou. 2020. "Estimating the Extent of Asymptomatic COVID-19 and Its Potential for Community Transmission: Systematic Review and Meta-Analysis". *Journal of the Association of Medical Microbiology and Infectious Disease Canada* 5(4):223–234.

Byrne, A. W., D. McEvoy, A. B. Collins, K. Hunt, M. Casey, A. Barber, *et al*. 2020. "Inferred Duration of Infectious Period of SARS-CoV-2: Rapid Scoping Review and Analysis of Available Evidence for Asymptomatic and Symptomatic COVID-19 Cases". *BMJ Open* 10(8):e039856.

Chatterjee, K., and V. S. Chauhan. 2020. "Epidemics, Quarantine and Mental Health". *Medical Journal, Armed Forces India* 76(2):125–127.

Chen, W., B. Zhang, C. Wang, W. An, S. K. Guruge, H.-K. Chui, *et al*. 2024. "A Metric of Societal Burden Based on Virus Succession to Determine Economic Losses and Health Benefits of China's Lockdown Policies: Model Development and Validation". *JMIR Public Health and Surveillance* 10:e48043.

Dönges, P., T. Götz, N. Kruchinina, T. Krüger, K. Niedzielewski, V. Priesemann, *et al*. 2024. "SIR Model for Households". *SIAM Journal on Applied Mathematics* 84(4):1460–1481.

Endo, A., M. Uchida, A. J. Kucharski, and S. Funk. 2019. "Fine-Scale Family Structure Shapes Influenza Transmission Risk in Households: Insights from Primary Schools in Matsumoto City, 2014/15". *PLoS Computational Biology* 15(12):e1007589.

Geffen, N., and M. Low. 2020. "Isolation of infected people and their contacts is likely to be effective against many short-term epidemics". *medRxiv preprint medRxiv: 2020.10.07.20207845*.

Giallonardo, V., G. Sampogna, V. Del Vecchio, M. Luciano, U. Albert, C. Carmassi, *et al*. 2020. "The Impact of Quarantine and Physical Distancing Following COVID-19 on Mental Health: Study Protocol of a Multicentric Italian Population Trial". *Frontiers in Psychiatry* 11(June):1-10.

Grijalva, C. G., N. Goeyvaerts, H. Verastegui, K. M. Edwards, A. I. Gil, C. F. Lanata, *et al*. 2015. "A Household-Based Study of Contact Networks Relevant for the Spread of Infectious Diseases in the Highlands of Peru". *PloS One* 10(3):e0118457.

Guo, Y., Z. Dou, N. Zhang, X. Liu, B. Su, Y. Li, *et al*. 2023. "Student Close Contact Behavior and COVID-19 Transmission in China's Classrooms". *PNAS Nexus* 2(5):pgad142.

House, T., H. Riley, L. Pellis, K. B. Pouwels, S. Bacon, A. Eidukas, *et al*. 2022. "Inferring Risks of Coronavirus Transmission from Community Household Data". *Statistical Methods in Medical Research* 31(9):1738–1756.

Jin, H., H. Wang, X. Li, W. Zheng, S. Ye, S. Zhang, *et al*. 2021. "Economic Burden of COVID-19, China, January-March, 2020: A Cost-of-Illness Study". *Bulletin of the World Health Organization* 99(2):112–124.

Liu, P., L. McQuarrie, Y. Song, and C. Colijn. 2021. "Modelling the Impact of Household Size Distribution on the Transmission Dynamics of COVID-19". *Journal of the Royal Society Interface* 18(177):20210036.

Liu, Y., A. A. Gayle, A. Wilder-Smith, and J. Rocklöv. 2020. "The Reproductive Number of COVID-19 Is Higher Compared to SARS Coronavirus". *Journal of travel medicine* 27(2):taaa021.

Møgelmose, S., L. Vijnck, F. Neven, K. Neels, P. Beutels, and N. Hens. 2023. "Population Age and Household Structures Shape Transmission Dynamics of Emerging Infectious Diseases: A Longitudinal Microsimulation Approach". *Journal of the Royal Society Interface* 20(209):20230087.

Mossong, J., N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, *et al*. 2008. "Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases". *PLOS Medicine* 5(3):e74.

Nash, D., S. Qasmieh, M. Robertson, M. Rane, R. Zimba, S. G. Kulkarni, A. Berry, W. You, C. Mirzayi, D. Westmoreland, A. Parcesepe, L. Waldron, S. Kochhar, A. R. Maroko, and C. Grov. 2022. "Household Factors and the Risk of Severe COVID-Like Illness Early in the U.S. Pandemic". *PLOS One* 17(7):e0271786.

Ponge, J., J. Suer, B. Hellingrath, and A. Karch. 2024. "A Standardized Framework for Modeling Non-Pharmaceutical Interventions in Individual-Based Infectious Disease Simulations". In *2024 Winter Simulation Conference (WSC),* 1106–1117. IEEE. https://doi.org/10.1109/WSC63780.2024.10838778.

Ponge, J., D. Horstkemper, B. Hellingrath, L. Bayer, W. Bock, and A. Karch. 2023. "Evaluating Parallelization Strategies for Large-Scale Individual-based Infectious Disease Simulations". In *2023 Winter Simulation Conference (WSC),* 1088–1099. IEEE. https://doi.org/10.1109/WSC60868.2023.10407633.

Potter, G. E., T. Smieszek, and K. Sailer. 2015. "Modeling Workplace Contact Networks: The Effects of Organizational Structure, Architecture, and Reporting Errors on Epidemic Predictions". *Network Science* 3(3):298–325.

Tseng, Y.-J., K. L. Olson, D. Bloch, and K. D. Mandl. 2023. "Smart Thermometer-Based Participatory Surveillance to Discern the Role of Children in Household Viral Transmission During the COVID-19 Pandemic". *JAMA Network Open* 6(6):e2316190.

Verity, R., L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, *et al*. 2020. "Estimates of the Severity of Coronavirus Disease 2019: A Model-Based Analysis". *The Lancet Infectious Diseases* 20(6):669–677.

Weigl, J. A. I., A.-K. Feddersen, and M. Stern. 2021. "Household Quarantine of Second Degree Contacts Is an Effective Non-Pharmaceutical Intervention to Prevent Tertiary Cases in the Current SARS-CoV Pandemic". *BMC Infectious Diseases* 21(1):1262.

Wu, Y., L. Kang, Z. Guo, J. Liu, M. Liu, and W. Liang. 2022. "Incubation Period of COVID-19 Caused by Unique SARS-CoV-2 Strains: A Systematic Review and Meta-Analysis". *JAMA Network Open* 5(8):e2228008.

Zhang, R., Y. Wang, Z. Lv, and S. Pei. 2022. "Evaluating the Impact of Stay-at-Home and Quarantine Measures on COVID-19 Spread". *BMC Infectious Diseases* 22(1):648.

AUTHOR BIOGRAPHIES

JOHANNES PONGE is a Research Assistant at the Chair for Information Systems and Supply Chain Management at the University of Münster, Germany. He holds an M.Sc. in Information Systems. His work focuses on the development of simulation-based decision support systems for infectious disease mitigation and intervention. He is a developer of the German Epidemic Micro-Simulation System (GEMS). His email address is johannes.ponge@ercis.uni-muenster.de.

JULIAN PATZNER is a Research Assistant at the Institute for Medical Epidemiology, Biometrics and Informatics at Martin Luther University Halle-Wittenberg, Germany. He holds an M.Sc. in Computer Science with a focus on Algorithm Engineering. His research interests include agent-based modeling, simulation, and optimization, particularly in the context of public transportation systems. His email address is julian.patzner@uk-halle.de.

BERND HELLINGRATH is a Professor and Head of the Chair for Information Systems and Supply Chain Management at the University of Münster, Germany. His research interests deal with the broader area of modeling and simulating with a distinct focus on the context of crisis-management and humanitarian logistics. His email address is bernd.hellingrath@ercis.uni-muenster.de.

ANDRÉ KARCH is a Professor of Clinical Epidemiology, and the Head of the Clinical Epidemiology Unit at the University of Münster, Germany. He holds an MD and an MSc in Epidemiology. His research in the field of infectious disease epidemiology focuses on dynamic transmission models for public health decision-making. His email address is akarch@uni-muenster.de.