WiFi mobility models for COVID-19 enable less burdensome and more localized interventions for university campuses

Vedant Das Swain[*] Jiajia Xie[*,†] Maanit Madan[*]
Sonia Sargolzaei[*] James Cai[‡] Munmun De Choudhury[*]
Gregory D. Abowd[*,◇] Lauren N. Steimle[†] B. Aditya Prakash[*,1]

[*]College of Computing, Georgia Institute of Technology
[†]H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology
[‡]Department of Computer Science, Brown University
[◇]College of Engineering, Northeastern University

March 16, 2021

Abstract

Infectious diseases, like COVID-19, pose serious challenges to university campuses, and they typically adopt closure as a non-pharmaceutical intervention to control spread early and ensure a gradual return to normalcy. These policies, like remote instruction (RI), reduce potential contact but also have broad side-effects on campus by hampering local economy, students' learning outcomes, and community wellbeing. In this paper, we demonstrate that university policymakers can mitigate these tradeoffs by leveraging anonymized data from their WiFi infrastructure to learn community mobility (WiMob) and in turn explore more granular policies like localized closures (LC). WiMob can construct contact networks that capture behavior in a variety of spaces, highlighting new potential transmission pathways and temporal variation in contact behavior. Additionally, WiMob enables us to design LC policies that close super-spreader locations on campus. On simulating disease spread with contact networks from WiMob, we find that LC maintains the same reduction in cumulative infections as RI while showing greater reduction in peak infections and internal transmission. Moreover, LC reduces campus burden by closing fewer locations, forcing fewer students into completely online schedules, and requiring no additional isolation. WiMob can empower universities to conceive and assess a variety of closure policies to prevent future outbreaks.

Keywords: COVID-19; mobility; modeling; public health; policy; non-pharmaceutical intervention; passive sensing; WiFi

[1]To whom correspondence should be addressed. Email: `badityap@cc.gatech.edu`

Introduction

University campuses are often hotspots for infectious disease outbreaks and hence are usually targeted for interventions. In the wake of the Coronavirus Disease (COVID-19) [41], the U.S. alone witnessed more than half a million cases at universities [54], and colleges are still left with decisions for operations in Fall 2021 [34, 44]. Controlling the disease at universities can be pivotal to securing the surrounding environment [5]. To reduce on-campus infections and the likelihood of superspreading events, a recommended form of non-pharmaceutical intervention (NPI) is partial closure of the campus [21]. During COVID-19, advancement in technology equips universities to continue operations by adopting a form of campus closure that relies on remote instruction (RI) [39]. As a consequence, the campus community has fewer opportunities to visit spaces to congregate and risk transmission, such as classrooms [1, 3]. One common approach campuses consider to design RI policies is to use enrollment data (En) to offer large classes online while offering other classes in person [7, 59]. In fact, during COVID-19, 44% colleges and universities in the U.S., primarily offered instruction online [52]. However, these policies can still have broad, negative, and indiscriminate impact on the community by forcing students into completely remote course schedules. Such policies can have adverse effect on learning outcomes [16], where students can lose close to 7 months of education [2]. Additionally, RI can disincentivize students to stay on campus and thus, incur losses in auxiliary revenue (e.g., boarding, parking, dining, etc.) [23, 15], with universities standing to lose up to $50 million because of unused services [63]. Even the local population unaffiliated with the university takes sustains losses to business due to university closures [28, 58]. Furthermore, with socioeconomic disparities and heterogeneous household contexts, the demands of remote instruction can lead to added anxiety and stress among students [11, 61]. Indeed, while campus closures have been deployed for prior epidemics as well, like influenza [9, 20], these policies have always been contentious because of significant disruption to educational and commercial activity, and social wellbeing [62]. Relying on RI, university campuses struggle to balance community health with the demands of learning, economy, and broad wellbeing [46]. Instead, there is a need for a more versatile approach to design closure policies that empowers policymakers to accurately assess impact of closure interventions and model more data-driven targeted intervention strategies.

This paper showcases a new approach to conceive closure policies on campus by leveraging data in their existing WiFi infrastructure. Our methodology, WiMob, involves constructing anonymized mobility networks of campus (Figure 1a), which helps determine extended periods of collocation — or "proximate contact" [27]— between individuals to describe contact networks on campus. Particularly, WiMob enables a more expressive toolkit for university policymakers that represents contact longitudinally and allows them to assess closure at the granularity of a room, suite, or hall. Thus, it lends to the design of targeted interventions that focus on localized closures (LC). We demonstrate the utility of WiMob with data collected over two years, of approximately $40,000$ anonymous occupants and visitors of the Georgia Institute of Technology (GT), a large urban campus in the U.S. — including about $16,000$ undergraduate students, $9,000$ graduate students, and $7,600$ staff members. In general, on comparing WiMob to En as an approach to model contact, we find that

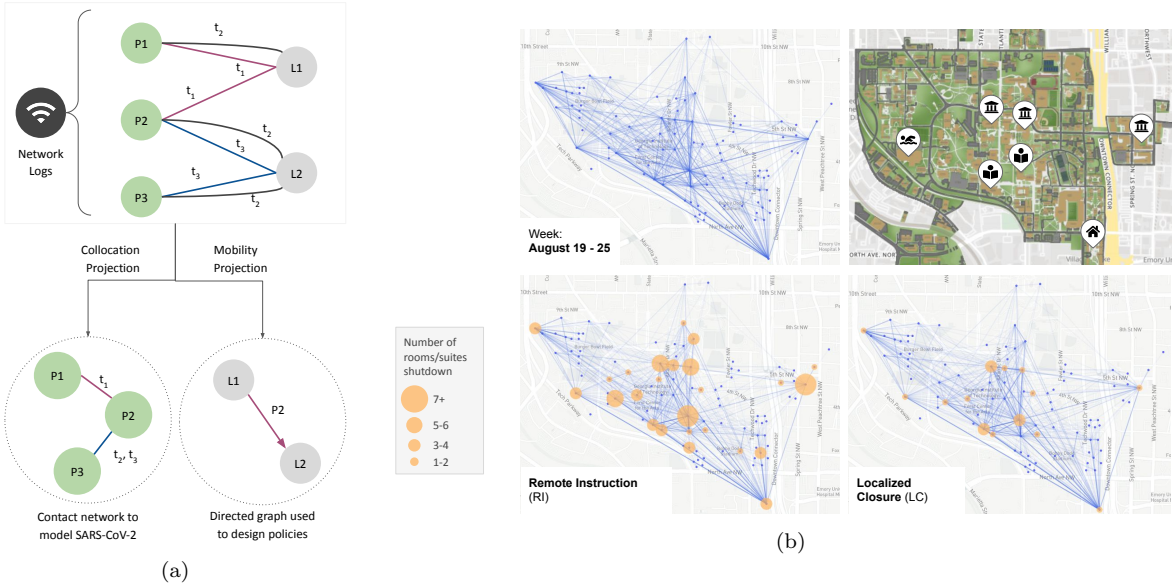

Figure 1: Anonymized WiFi network logs help model mobility on campus (WiMob) and target spaces for localized closures (LC) (a) WiFi network logs reflect timestamps when people's devices associate with different access points (APs) on campus. WiMob mines these logs to characterize mobility as a bipartite graph that describes people (e.g., $P1$, $P2$) visiting locations (e.g., $L1$, $L2$) on campus during different times (e.g., t_1, t_2). Since people's devices can proxy their presence, we estimate collocation (e.g., $P1$ and $P2$ were collocated at $L1$ at t_1), and movement ($P2$ dwelled at $L1$ and then at $L2$). The collocation graph forms the basis of the contact structure for our ABM. (b) WiMob highlights mobility behavior to evaluate and inform policy. (b)–*top-left:* Mobility on campus between the top 100 most frequented locations on the GT campus in the Fall semester of 2019. Edges only connect points of significant dwelling and thus do not represent pedestrian routes. (b)–*top-right:* For reference, we annotate one large on-campus housing facility (south-east), the recreation center (west), two learning commons (center), and three highly frequented academic buildings. Some of the highest centrality locations are in these buildings. (b)–*bottom-left:* RI is a form of broad closure which affects a large number of students and locations.(b)–*bottom-right:* By contrast, we propose to use WiMob to parsimoniously identify a small set of spreader locations within buildings and design LC policies. With these policies, we aim to provide equal or better control on the disease spread, and yet minimize the burden on campus compared to RI.

WiMob captures contact behavior at a community scale for a variety of campus spaces, describes temporal variations in contact, and provides a better estimate of local context by being aware of occupancy and the non-student population. Leveraging WiMob also reveals that En overestimates the impact of RI on reducing contact on campus. Hence, we propose a less burdensome alternative to RI, by deriving more targeted LC policies based on WiMob (Figure 1) (indeed En is too coarse-grained for designing targeted LC policies).

We further exhibit that LC presents better disease control outcomes than RI by constructing and simulating an agent-based model (ABM) over the WiMob contact networks, calibrated with GT on-campus COVID-19 cases from the Fall semester of 2020 [25] and infection rates from Fulton County [40]. To compare the effect of interventions, we describe a counterfactual semester that is unaltered by other policy–induced behaviors of 2020 by leveraging WiFi data from Fall 2019 to determine the contact structure of the simulation.

This model helps assess the effectiveness of closure NPIs (Figure 1b) under various behavioral scenarios. We find LC is comparable to RI in controlling total infections but more effective at reducing the peak infections and internal transmission. Additionally, LC targets fewer locations, forces fewer students into fully online schedules, and does not isolate any more people than RI – illustrating that WiMob can help universities devise highly-specific closure policies, like LC, which can contain disease spread and mitigate campus disruption in comparison to RI policies.

Our methodology also promises other advantages. Mobility generally has been used to dynamically model disease spread of influenza [47], rubella [60] and COVID-19 [3, 45]. For the latter, various studies show the effectiveness of mobility restrictions at a regional–, or city–level [66, 10, 6, 37, 30]. Previous studies that use mobility information to model disease spread and interventions typically rely on cell tower localization or aggregating GPS information from mobile phones. Neither of these data sources is easy to access for university campuses. In the past, studies to infer campus mobility have relied on accessing user devices with specialized data logging applications [12, 17, 49], but these approaches are typically constrained for disease modeling because they require mass adoption to represent the entire community and continuous maintenance of software is needed to capture longitudinal behavior changes. In contrast, this study repurposes already existing managed WiFi networks to model mobility, which provides room level granularity for mobility [18, 57, 13, 55] and consequently indicates proximate contact [27]. Much like En, universities internally archive such data over a long term for other purposes and do not need to install any additional surveillance infrastructure to access it. With the appropriate privacy considerations, a university can obtain this data at a low cost, continuously and unobtrusively.

The possibility of pandemic still looms large in the future [32, 24]. As campuses prepare for the upcoming Fall semester and unforeseen contagious diseases of tomorrow, WiMob presents an attractive and practical method to inform better public health policies.

Results

We present two sets of analyses in our work. The first set contrasts structural characteristics of contact networks described by WiMob with current practices that use En. In the next set, we use WiMob to model both RI and LC interventions and analyze their differences in terms of dynamic disease-control outcomes and burdens to campus.

WiMob provides local, holistic and dynamic structural insights for contact networks on campus

Studies on RI policies tend to assume that contact in universities is largely informed by En—transcripts showing which courses a student is registered for. En can provide structural insights on density of connections and disease transmission paths to inform modeling disease simulations [26]. However, such static data can overestimate attendance and ignore overlap

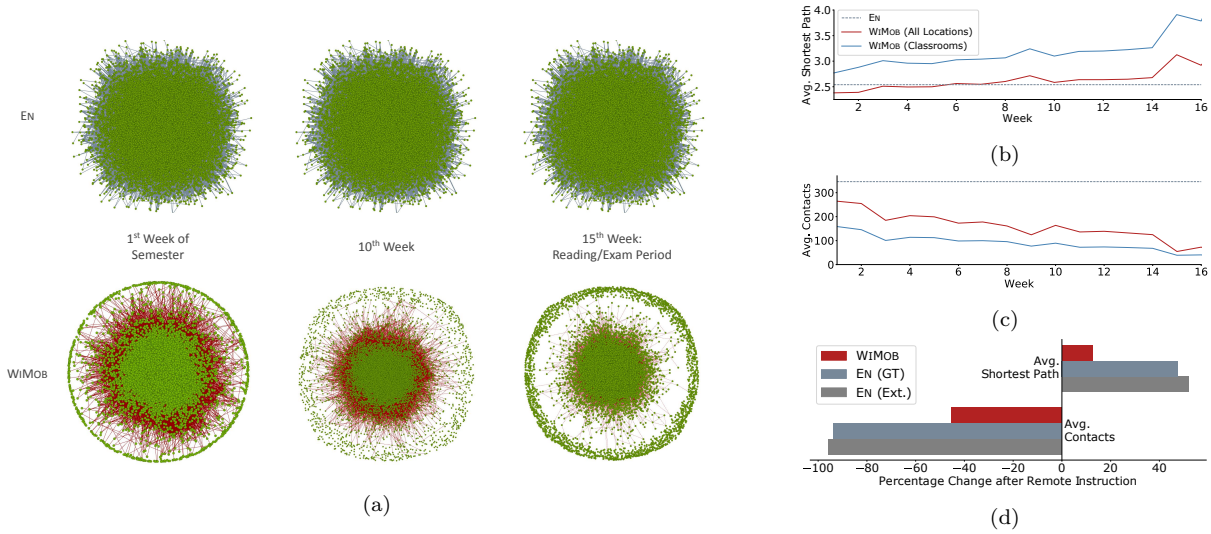

Figure 2: Results show difference in structural characteristics of contact networks from En and WiMob. (a) En describes contact based on classes that students are expected to attend based on their enrollment. En assumes 90% of students to be connected in a single component, but WiMob reveals that on given week only 69% of the academic population is in the largest component (those that do not visit campus are isolated and shown in the circumference). In the first week of Fall 2019, both En and WiMob show high connectivity. However, En does not change, but WiMob reveals that density of connections changes over the semester. (b, c) The difference in the two disease models is because of the difference in structural characteristics of networks built with WiMob compared to En. (b) En depicts campus contacts to be connected closely into a "small world". WiMob shows that contacts evolve over time. As mobility captures interactions outside classrooms we observe that for the first 6 weeks the shortest transmission path between people is shorter than what is reported by En. (c) Enrolling into a course does not necessitate physically collocating with the class for extended periods (students can also choose to be entirely absent). WiMob reflects this behavior and highlights a decline in average contacts over time. (d) These structural differences can help policymakers anticipate the effect of closure policies by describing how it fragments the underlying contact network. En shows that remote instruction leads to a 94% reduction in contacts and 50% increase in transmission path length (similar to numbers reported in [59], shown as En (Ext.)). However, the estimate is significantly lower with WiMob. As a result, WiMob emphasizes the limits of remote instruction policies and in turn motivates new policies that can be designed and evaluated with actual on-campus behavior.

between courses (via instructors) and organic interactions outside classes (e.g., waiting areas, dining, parties, and extra-curricular activities). Therefore, using En can overemphasize the disease-mitigating structural changes to the network by RI interventions. By contrast, WiMob is more grounded in community behavior as it captures multiple scheduled and serendipitous contact situations dynamically over the semester. We compare the features of contact networks constructed with WiMob, against networks constructed with En using data from GT for Fall semester of 2019 (August 19 – December 14), which is before any COVID-19 reported cases in the U.S. En approximates contact based on students enrolling for classes that could potentially collocate them in the same room during lectures. WiMob infers contact when any two individuals actually collocate near the same WiFi access point [13, 55] for extended period (see explanation in SI WiFi Mobility). We find that WiMob renders new insight into contact on campus that is invisible to the En methodology.

WiMob characterizes temporal variation in proximity

Variation in contact over the semester would naturally impact the severity of disease spread. However, En describes a static network that does not capture such dynamics (Figure 2a). Instead, we find that WiMob shows contacts get sparser as the semester progresses. Figure 2c presents a notable decline in contacts after the first two weeks, which coincides with multiple orientation seminars and the so-called "course shopping" period of Fall 2019. In fact, contact decreases considerably in classrooms, with a steeper slope possibly because of reduction in attendance. WiMob is able to reveal other observable changes, such as drop in contacts during exam period (week 15) and increase after fall recess (week 10). Since, En renders a highly connected static network, which can miscalculate the speed at which a disease spreads. By contrast, the longitudinal behavior represented by WiMob can help universities anticipate disease spread more accurately.

En overestimates contact-based risk

Campuses can assess risk of an outbreak by characterizing the number of individuals that would be at risk of infection through contact. In fact, En indicates that 99% of the individuals on campus are clustered in a single component — if any of them would have been infected in Fall 2019, everyone in the component would be at risk. From the lens of En a virus can exhaust an entire population with infection very early. However, WiMob shows that only 69% of the population is connected in a single component (Table S3). This difference is because WiMob can distinguish how many individuals are active on campus. Therefore, WiMob provides a pragmatic estimate of risk by grounding it in local occupancy and helps campuses budget for resources better.

WiMob reveals different paths for disease transmission

Reports suggest that a key contributor to cases in the pandemic is actually clustering of individuals in non-academic spaces [39]. However, En does not depict a holistic view of campus contact. It is limited to classrooms and, therefore, fixates on contacts in lectures, while ignoring other spaces. In fact, WiMob shows that in the first 6 weeks of Fall 2019, the shortest path among individuals is smaller than that approximated by En (Figure 2b). With WiMob, we observe new paths in the contact network from situations outside classes. On a given week, WiMob shows the average shortest path with contact is $3.26(\pm 0.5)$ when only considering lectures, whereas capturing all contexts reduces the average shortest path to $2.67(\pm 0.28)$. Characterizing shorter pathways is crucial for policymakers as closure policies by design aim to disconnect these pathways.

En overemphasizes the impact of remote instruction

Prior work uses En to posit that RI reduces contact and in turn significantly fragments the network for disease spread in universities [59, 7]. To compare policy effectiveness with WiMob, we operationalize RI in our study:

Remote Instruction (RI) The status quo for data-driven policies offers strictly online instruction for large class enrollment, while continuing the other classes in person. For En we implement this by removing connections between students who are only in contact through courses where size ≥ 30. For WiMob we remove connections between students if they are only connected because of collocations during scheduled lectures of such courses.

We evaluate the effectiveness of such a policy if it were applied in Fall 2019, with both WiMob and En. Figure 2d shows that RI with En reduces contact by 94% and increases shortest path by 50%. However, the same intervention with WiMob shows a relatively milder impact (contact reduction 45%; shortest path increase 11%). This reinforces that contact outside courses are significant and remain unaffected by enrollment-oriented policies like RI. WiMob provides a more encompassing view of the structural effects to a network and motivates design of more impactful closure policies.

LC yields better infection reduction outcomes with lower burden

As outlined above, En does not capture comprehensively the contact on campus. A campus is composed of many different spaces and En does not have the flexibility to design closure of such spaces or assess its impact. These drawbacks naturally motivate a new approach to design interventions. Since WiMob mitigates the limitations of En, we leverage it to demonstrate the effectiveness of localized closure or LC.

We evaluate the community health outcomes and burdens to campus of closure interventions by simulating COVID-19 with our ABM that uses WiMob to define the contact structure for each day. This is overlayed by a modified SEIR compartmental model for COVID-19. GT also had implemented a robust surveillance program on campus. Hence we calibrate the ABM on the positivity rate for COVID-19 for GT [25] in the first 5 weeks of Fall 2020 also incorporating external seeding from the surrounding Fulton County, GA [40]. We validate our model by predicting future trends for the rest of Fall 2020. For robustness, we perform additional calibrations by varying time windows and university context (details in SI Sensitivity Analyses). We study interventions by applying the ABM over the contact networks produced by WiMob with data from Fall 2019 — a counterfactual to Fall 2020 if no closure had occurred (see SI Simulation Model for further details).

WiMob can model RI and LC interventions with various configurations

In addition to RI, we model LC, which we formalize:

Localized Closure (LC) Prior works show a few locations are responsible for majority spread [10] and restricting movement between regions leads to greater control [33]. We intuitively identify rooms–level spaces that are highly central location nodes in the network. We remove contacts between people who are only connected because of collocating at these locations. While, we employ various centrality algorithms to identify

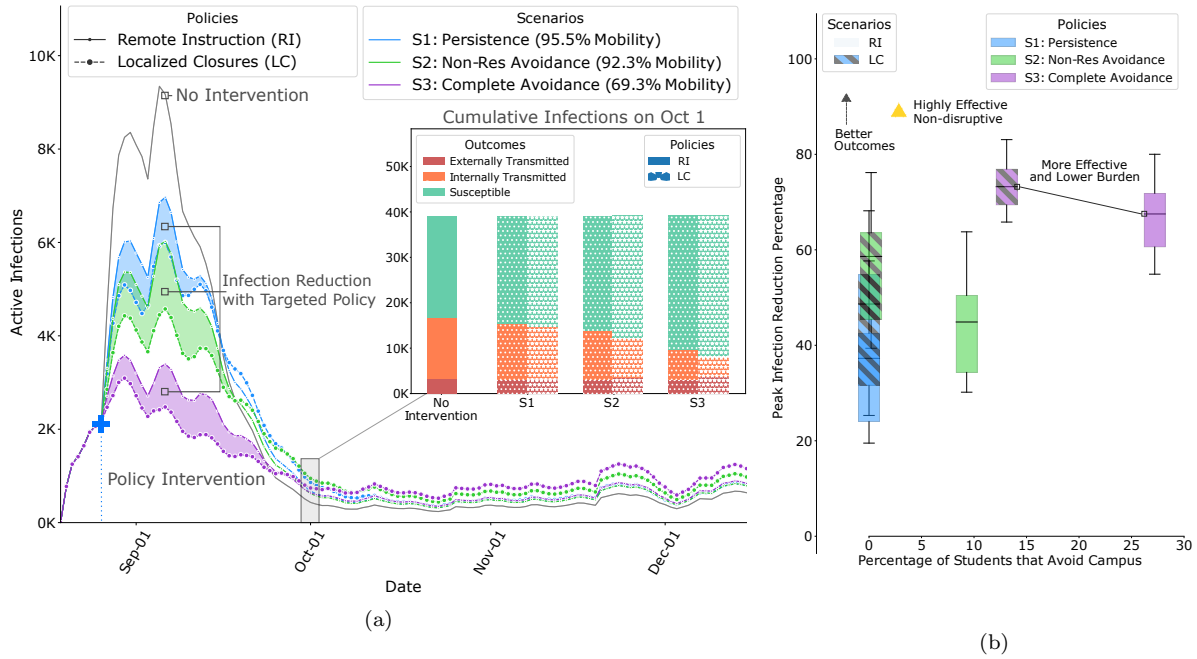

Figure 3: Results of policy interventions with our calibrated ABM on contact networks from Fall 2019, derived from WiMob (*a*) LC shows improved outcomes (shaded regions) even when constrained to the same restrictions of broader shutdown policies. This holds true especially during the first month after the policy is implemented and is consistent regardless of the behavior scenarios. After the first month, the active infections decline for non-intervention and broader shutdowns appear to be lower, but the cumulative infection count is still higher (inset). The overall decline in infections could be a function of reduced contact (Figure 2b) as well as fewer external infections from the county. It is important to note that new infections are a proportion of susceptible individuals in the system. Since LC restrict spread early, it leaves more people susceptible (to both internal and external infections) later in the semester. (*b*) LC not only provided better outcomes (for peak and internal infection reduction) they also tend to be less burdening in terms of certain cost dimensions. For instance, in certain scenarios LC forces fewer students into fully online schedules and therefore keep more students on campus.

such locations, but for the results discussed in this section we use *PageRank* [42]). Details in SI Identifying Locations for Closure.

We find that, if COVID-19 spread through Fall 2019 (a regular semester), the cases rise after 7 days (Figure 3a). Therefore, we apply both RI and LC interventions after the first week.

To make the comparisons between the closure policies, we establish fixed budgets to design LC based on the resource utilization on RI We consider 2 kinds of budgets, (i) mobility reduction — to depict space use on campus, and (ii) risk of exposure — to reflect testing capacity. Also note, response to closure policies can lead to unpredictable side-effects in campus behavior, particularly when a student's schedule is entirely online. Therefore, we design policies within three behavioral scenarios (each with a varying budget):

- *S1: Persistence*: Irrespective of the locations closed or classes restricted, individuals continue their other visiting behaviors.

- *S2: Non-Residential Avoidance*: Non-residential students stop all visits to campus if they enrolled in at least 3 courses and the policy forces their entire academic schedule online.
- *S3: Complete Avoidance*: Both residential and non-residential students avoid campus if they have at least 3 courses and all move online.

WiMob facilitates the conception of expressive closure interventions. To devise interventions,WiMob estimates how RI uses the budget and then designs LC to match this budget under every scenario. Table 1 describes how the budget for each policy varies. Additional details are present in SI Modeling Policy and Scenarios.

We present differences between LC and RI based on three infection reduction outcomes; peak infections (maximum active cases on a given day), internal transmission (exposure from infected individuals on campus), and total infections (cumulative cases at the end of the semester). Additionally, we measure the burden of policy interventions with the number of locations closed — requires resources to monitor and maintain super-spreader locations, the percentage of students that avoid campus — disruption to learning outcomes [16, 11], and the percentage of individuals completely isolated — worsens mental wellbeing [48].

LC cause greater reduction in peak infections, while affecting fewer locations

Controlling peak infections relaxes the burden on a university to support positive cases for any given day, and allows resources to be distributed over time. In all scenarios, of our simulation of Fall 2019, we observe that the peak reduction is significantly better in LC (Figure 3) than RI. While RI impacts 58 different locations (classrooms and lecture halls), in *S1* and *S2*, LC achieves better outcomes by closing fewer locations. For example, in *S2*, RI achieves a 28.9% peak reduction, but LC shows a reductions of 49.3% (mobility budget) and 48.1% (exposure risk budget). This is attained by closing 38 or 50 locations respectively. Therefore, with such policies, policymakers need to restrict fewer locations to remarkably minimize the pressure of active infections on campus (e.g., diagnoses, treatment, quarantining).

LC lead to comparable reduction in total infections, while keeping more students on campus

Universities want to minimize the number of infected cases while ensuring majority of the population remains active on campus to continue successful operation. The total number of infections reduced by both LC and RI are similar. While the differences between policies are statistically significant (Table S3) in some scenarios, the magnitude of these differences might not be practically as important. In contrast, the impact the policies have on the student schedules is remarkably different. RI forces multiple students to adapt to fully online schedules. In the Scenario *S2*, 9% of students do not visit campus and in *S3*, 27% of students do not visit campus. On the other hand in LC the number of students expected to avoid campus can be as low as 0 and never exceeds more than 12%. Besides sustaining economic

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	RI	LC		RI	LC		RI	LC	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	25.34(±12)	36.92(±14)**	34.30(±13)**	35.44(±10)	49.33(±11)**	52.19(±10)**	61.62(±7)	69.34(±5)**	64.44(±6)**
Total Infections (%)	6.99(±5)	10.63(±6)**	8.19(±5)**	14.88(±4)	13.96(±6)*	15.67(±6)	33.00(±5)	33.4(±5)	26.94(±5)**
Internal Transmissions (%)	17.13(±9)	22.62(±11)**	21.01(±11)**	27.58(±8)	35.35(±12)**	39.20(±11)**	54.00(±8)	70.89(±7)**	60.90(±9)**
Burdens on Campus									
Locations Affected	58	18	19	58	38	50	58	192	124
Students Avoiding (%)	0	0	0	9.30	0.20	0.45	27.21	12.45	6.57
Completely Isolated on Campus (%) 5.42	8.40	8.40	5.95	5.72	5.71	7.09	5.18	5.23	

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of LC with RI. We find that LC leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes are comparable in general but can vary by specific scenarios. In addition, every policy also exerts some burden on campus, either in terms of locations affected, students avoiding campus or isolation. We observe that LC policies focus on fewer locations (except in *S3*). Moreover, these policies affect fewer student's schedules and therefore fewer people avoid campus due to completely remote schedules. Finally, LC does not increase the percentage of people completely isolated on campus (p-value: < 0.01:*, < 0.001:**).

Table 1: Comparison of different policies in terms of controlling the disease and impacts on campus in Fall 2019.

loss to the campus, remote instruction can increase anxiety among students and hinder learning outcomes [11, 61]. Compared to RI, LC offers policymakers a way to defend against turnover in the student population, without compromising overall control of disease spread (Table 1). Limiting the number of students that avoid campus helps preserve on-campus businesses [28, 58] and minimally disrupts the student wellbeing.

LC cause greater reduction in internal transmission without causing further isolation on campus

Universities are responsible for limiting spread on campus, but they must also ensure that aggressive policies do not worsen mental wellbeing of the community. In terms of internal transmission the reduction is significantly larger with LC (Table 1). However, when LC restricts the infections early in Fall 2019, it leaves more individuals susceptible to external transmission. College student behavior outside campus on weekends and breaks is known to impact local transmission [14]. When policymakers consider LC they should also consider policies on re-entry or required testing based on off-campus activities. In terms of isolating individuals on campus, it's notable that LC and RI are similar in *S2*. Interestingly, in *S3*, where LC closes more than 100 locations, the percentage of isolated individuals per week is less than that of RI. This finding implies that LC can keep individuals on campus without forcing them into complete isolation. Here "isolation" refers to no form of proximate contact with any individual on campus — extreme social distancing where individuals are not even collocated in the same suite or hall. While social distancing is a recommended coun-

termeasure for COVID-19 [1], complete isolation can have adverse effects on psychological wellbeing [48, 36, 43]. Staying completely isolated on campus can increase loneliness and limit social connectedness [36], which are both related to depression [48]. Although the proportions are similar (Table 1), LC does not necessarily isolate the same sets of individuals. This qualitative difference could also explain the difference in internal transmissions — LC could be isolating individuals who are less likely to spread the virus.

LC identifies a wider variety of auxiliary spaces

By using WiMob to design LC we are able to identify locations for closure at the granularity level of rooms, including unbound spaces such as lobbies and work areas. First, in *S1*, we find that most locations that LC targets are a subset of the auditoriums–like rooms where large classes would take place in Fall 2019. Note, LC needs to restrict only a few such spaces to be under the same budget as RI. This is because, under *S1*, RI policies only alter visits to lectures, while these spaces are used for other purposes during other times (e.g., club activities and seminars). We also note that LC targets 'high traffic' locations like conference center lobbies which are typically used as waiting areas or for networking events. Next, in Scenario *S2*, we see that in addition to spaces mentioned earlier, interestingly LC further restricts the use of smaller rooms (occupancy $13-35$) which would not be affected by RI (as only classes of size ≥ 30 are offered online). LC also targets areas in the recreation center (which includes locker rooms and indoor courts for $4-20$ people). This insight indicates that our methodology WiMob is sensitive to other student activities. Moreover, we also find a selection of spaces that would not be frequented by the undergraduate population, such as lab areas and facility buildings like the police station. Lastly, in Scenario *S3*, LC targets closure of activity in far more spaces than RI. However, the better outcomes can be attributed to the fact that LC diversifies the potential restriction areas. LC now restricts heavily used small study rooms or breakout rooms (for $1-6$ people). Furthermore, it restricts use of spaces where multiple small groups of people can organically assemble, such as cafes, dining halls, and reading areas. We also observe that LC restricts activity in about 10 Greek Houses but does not target other housing areas — demonstrating its ability to restrict social behavior that could amplify disease spread. Figure S22 shows the diversity in locations for various LC policies.

Sensitivity and robustness analyses

The results above use an ABM calibrated on the positivity rate of the first 5 weeks of Fall 2020. At a university, this rate can be influenced by many latent factors (e.g., mask-wearing, hand washing, distancing, and compliance). To study any effect of these variations on our results, we also calibrated on different time windows throughout the semester. We calibrate on weeks $5-9$ and $10-14$ in Fall 2020, and validate on the remaining semester. In both cases, compared to RI, we find that LC still exhibits better reduction in peak infections (up to 90%) and internal transmission (up to 77%). In the original calibration, LC also significantly reduces total infections and maintains the same level as RI, but with the new periods we find

total infections are substantially less than RI (Table S9 and Table S10).

Another important variable for positivity is the wider context of the campus e.g. urban/rural, the surrounding county, city, etc. To investigate this, we also calibrated our ABM on the positivity rate of different universities in the US in Fall 2020 (along with information from their county to seed external cases). Consider this as a hypothetical where the mobility of the GT community remains the same but disease outcomes resemble a different campus. We calibrate on data from University of Illinois at Urbana-Champaign and University of California, Berkeley. We find no remarkable differences from our findings with GT (Table S11 and Table S12).

Discussion

When facing a pandemic, non-pharmaceutical interventions (NPI) are the first line of defense for universities to respond to contagious diseases like COVID-19 [19, 38]. On a campus, a common form of NPI is closure [29]. Universities consider enrollment data (En) to design remote instruction (RI) for closure to support continued operations safely [59]. However, En can misconstrue contact on campus and RI policies can have broad impacts despite their effects on curbing the disease spread. This paper demonstrates that repurposing logs from a managed WiFi network (WiMob) can help design effective localized closure policies (LC). We show that WiMob uncovers rich contact dynamics and provides policymakers multiple dimensions to design policies like LC. We simulate COVID-19 with an ABM that harnesses WiMob to compare RI and LC. Our results present evidence that LC can lead to improved infection reduction outcomes, while simultaneously relaxing burdens on the campus community caused by coarse-grained broad RI policies.

Generalizability for Other Contexts

In practice LC policies should be deployed on campus in conjunction with the other tools as well like testing, tracing, and quarantining. WiMob can complement disease-specific knowledge to identify closure spaces. For example, small indoor spaces with poor ventilation increase the risk of infection for COVID-19 [51], while other algorithm-identified locations for closure might not require closure because mask-wearing and testing have high compliance among users of that space. Further, as a pandemic progresses and public health guidance develops [50], with WiMob, campuses can regulate the restriction of LC policies and anticipate the path to 'normal' operations [34, 44] Moreover, WiMob captures various spillover effects that cannot be captured in methods like En. For instance, with WiMob we observe that the mobility in Fall 2020 was 39% of that in Fall 2019 because the on-ground policies lead to certain staff working remotely as well. With additional information, WiMob enables policymakers to model such scenarios and design alternatives like LC with new budgets. Policymakers in universities can use WiMob as a versatile tool to explore dynamic intervention strategies as well. In this paper, LC interventions are non-adaptive fixed policies throughout the semester. Since staggering policy restrictions could have variable impact on campus [65].

However, WiMob could help identify locations from different mobility phases (e.g., the same week from a prior semester) and assess the effects of closure policies. Additionally, depending on campus priorities and resource limitations, different campuses can use this same data to model policies differently. The effectiveness of reopening policies is expected to be sensitive to a campus' specific context that includes physical infrastructure, overarching guidelines, and human compliance [5]. For certain campuses policies might not need to be constrained by exposure risk as testing might be frequent, ubiquitous, and voluminous. Other campuses could have limits on quarantining capacity. Policymakers might even consider the cost trade-offs by actually forecasting actual financial losses incurred by reduction in mobility [6], or valuate loss of services based on community needs [53]. We elaborate on these considerations in the SI Implications for Policy Design.

Operational Considerations

Beyond assessing cost-benefits, universities also need to consider practical methods of obtaining, storing, and processing mobility of the community as WiMob. University can access logs from the managed network internally as it is passively collected. Moreover, it does not require any new form of surveillance sensing but universities must revise terms of use and stay sensitive to community perspectives. While aggregate data on population mobility is valuable for many applications [64], which includes informing pandemic response [8], the major privacy challenge with localization data is to avoid accumulation [56]. Instead, operational applications need to conceive approaches that only retain processed insights on locations to shutdown but not individual data. Similarly, any operational use needs to have pre-established access limitations on what stakeholders can learn from the data [4] (e.g., decision-makers can only get a list of candidate locations to close). In the SI Discussion, we further detail approaches to reconcile privacy, ethics and legal considerations.

Limitations and Future Work

Lastly, for future investigations of better closure policies, researchers and policymakers need to be cognizant of the limitations of our work. Our analyses do not represent heterogeneity among individuals and therefore our simulation does not account for intrinsic vulnerabilities [31, 43, 22] and difference in mobility behaviors of demographic groups [10]. WiMob can be extended with other streams of data to introduce variability in the population and devise new forms of LC to protect the most vulnerable community members. Additionally, our work explores the extremes of the range of behavioral responses to closure interventions. Henceforth, researchers and policymakers can model more nuanced spillover effects that interpolate between the scenarios we describe, as well as extend them. Further discussion in SI Limitations and Future Work.

Methods

This section summarizes (i) the data used to derive contact networks and policies, and (ii) the dynamics of our simulation model and calibration approach. Additional information for every subsection is present in SI Methods.

WiFi Mobility

Data Use and Access

The IT management facility at GT accumulates WiFi access point logs over time. This is common in most universities with managed WiFi infrastructure. We actively collaborated with IT management to define safety and security safeguards that allow us to obtain a de-identified version of these raw logs. Before accessing the data we established a data-use agreement and an ethics protocol that was approved by the Institutional Review Board (IRB). For the WiFi data, we were provided access to logs from Fall 2019 and Fall 2020. We process these logs to characterize mobility (WiMob) and it encompasses all $40,000$ unique individuals that connected to the network via $6,959$ different access points [13]. The logs do not contain any personally identifiable information and locations are also coded. For En we only use aggregate insights for enrollment, which are derived from course registration transcripts. Note, we do not cross-identify any students. We use publicly accessible course schedules to approximate schedules of de-identified nodes and infer if they are students or staff, and non-residential or residential. We elaborate on our data in SI Data.

Contact and Movement Networks

WiMob leverages the logs to create bipartite graphs K_t, for each day t, which connect P users to L access point locations (Figure 1a). Any edge, $\{p, l\}_i$ indicates the ith instance when a p was dwelling at l. These edges describe the time period of dwelling. Subsequently, by comparing all edges in K_t we can infer if different individuals are collocated near an AP to create a contact network, G_t, for each day t — between collocated $p \in P$. These networks feed into the ABM at every time-step. Similarly, by inspecting the sequence of dwelling locations for any p in graph K, we compute a mobility network, H_t — between locations $l \in L$. We provide more details of our approach in SI Data Processing and in SI Modeling Collocation and Movement.

Modeling Policies

We compare the disease outcomes and burdens of 2 policies, Remote Instruction (RI) and Localized Closure (LC), both of which are modeled with WiMob. For RI we infer enrollment size of each course in Fall 2019 by determining the number of unique individuals that visit lecture locations during scheduled times. After the first week, we apply the RI by removing all visiting edges in K_t for any $l_c \in L_{\mathsf{RI}}$ if visits were during lecture times of course c with an enrollment ≥ 30. This helps create counterfactual contact networks G'_t. The removal of

edges from K describes the mobility budget of RI and the structure of G'_t indicates the risk of exposure budget. We design LC with these budgets by identifying locations for closure (L_{LC}) with different algorithms, such as *PageRank, Eigenvector Centrality, Load Centrality, and Betweenness Centrality*. When a location is closed, we remove all edges in K_t connected to any $l_x \in L_{\mathsf{LC}}$. We aggregate the movement graph H_t over a week and apply the algorithms to identify locations. Subsequently, we identify the number of top-ranked locations to remove such that the resultant counterfactual contact network $G"_t$ has is within 1% of the budget. The budgets vary for different behavioral scenarios and we only compare policies within the same scenario. This is further elaborated in SI Modeling Policies and Scenarios.

Disease Simulation

Agent-Based Model

We construct an agent-based model (ABM) that captures the spread of COVID-19 between individuals active on campus. This ABM leverages the contact networks produced by WiMob. The simulation iterates a time-step each day and the underlying contact networks i.e., G_t for no interventions, G'_t for RI, and $G"_t$ for LC. Our ABM follows a modified version of *susceptible–exposed–infectious–removed* (SEIR) template that disambiguates the *infectious* compartment into *asymptomatic* and *symptomatic*. New infections are introduced to the model either externally or internally. External transmission arises because individuals can contract the virus outside campus and bring the infection back for local spread [24, 39]. We adopt data of positive cases from Fulton county [40] with a scaling factor α to estimate the probability that a *susceptible* individual, who is active on campus, was infected from interactions outside campus. This is to account for any commute outside campus during the pandemic. Internal transmissions are determined by p, as the probability of *susceptible* individuals in contact with an *infectious* one. We calibrate the parameters related to disease transmission by training and validating our models on the positivity rate reported by GT surveillance testing [25]. SI Agent-Based Model details the disease progression and describes the various parameters.

Calibration

We estimate the ranges of optimal parameters for disease transmission by minimizing the root means square error (r.m.s.e) between the Georgia Tech surveillance testing positive rates [33, 25] and the observed positivity rate of the model every week— percentage of new *asymptomatic* out of the total testable population. The surveillance testing conducted by Georgia Tech is designed for detecting individuals who contracted Covid-19 without showing Flu-like symptoms within the community [33]. We calibrate the model on the positivity rates on the first 5 weeks of Fall 2020. To attain a point estimation of the optimal parameters, we fit the model to predict trends in the remaining weeks by running a numerical optimization algorithm, Nelder-Mead [25]. To account for quantitative uncertainty, we estimate a range of parameters, within 40% of optimum r.m.s.e. Note, this calibration

characterizes latent factors associated with pandemic-related cautious behaviors, including the relationship with external transmission. And these factors could be related to "county characteristics, partisanship, media consumption, and racial and ethnic composition" [1]. Since the effectiveness of shutdown policies can vary by time period and county. In SI Sensitivity we discuss hypothetical variations where the mobility behavior of GT remains constant but disease outcomes change based on time period of calibration and positivity rates from universities at different counties in the U.S. See SI Calibration for details on the calibration process and results are in Table S4.

Acknowledgements

This paper is based on work partially supported by the NSF (Expeditions CCF-1918770, CAREER IIS-2028586, RAPID IIS-2027862, RAPID IIS-2027689, Medium IIS-1955883, NRT DGE-1545362, CCF-2115126), CDC MInD program, ORNL, and Semiconductor Research Corporation (in collaboration with Intel Labs). Some research personnel were supported by internal seed funding from the Georgia Institute of Technology and Georgia Tech Research Institute. Other computing resources were provided by the Office of Information Technology at Georgia Tech. The authors thank Di Wu, Hanna Hamilton, and Dima Nazzal (Georgia Institute of Technology) for their analysis of En.

References

[1] M. Andersen. Early evidence on social distancing in response to covid-19 in the united states. *Available at SSRN 3569368*, 2020.

[2] J. P. Azevedo, A. Hasan, D. Goldemberg, S. A. Iqbal, and K. Geven. *Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes: A set of global estimates*. The World Bank, 2020.

[3] H. S. Badr, H. Du, M. Marshall, E. Dong, M. M. Squire, and L. M. Gardner. Association between mobility patterns and covid-19 transmission in the usa: a mathematical modelling study. *The Lancet Infectious Diseases*, 20(11):1247–1254, 2020.

[4] E. Bagdasaryan, G. Berlstein, J. Waterman, E. Birrell, N. Foster, F. B. Schneider, and D. Estrin. Ancile: Enhancing privacy for ubiquitous computing with use-based privacy. In *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society*, pages 111–124, 2019.

[5] J. C. Benneyan, C. Gehrke, I. Ilies, and N. Nehls. Potential community and campus covid-19 outcomes under university and college reopening scenarios. *medRxiv*, 2020.

[6] S. G. Benzell, A. Collis, and C. Nicolaides. Rationing social contact during the covid-19 pandemic: Transmission risk and social benefits of us locations. *Proceedings of the National Academy of Sciences*, 117(26):14642–14644, 2020.

[7] M. Borowiak, F. Ning, J. Pei, S. Zhao, H.-R. Tung, and R. Durrett. Controlling the spread of covid-19 on college campuses. *arXiv preprint arXiv:2008.07293*, 2020.

[8] C. O. Buckee, S. Balsari, J. Chan, M. Crosas, F. Dominici, U. Gasser, Y. H. Grad, B. Grenfell, M. E. Halloran, M. U. Kraemer, et al. Aggregated mobility data could help fight covid-19. *Science (New York, NY)*, 368(6487):145–146, 2020.

[9] S. Cauchemez, A.-J. Valleron, P.-Y. Boelle, A. Flahault, and N. M. Ferguson. Estimating the impact of school closure on influenza transmission from sentinel data. *Nature*, 452(7188):750–754, 2008.

[10] S. Chang, E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky, and J. Leskovec. Mobility network models of covid-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, 2021.

[11] I. Chirikov, K. M. Soria, B. Horgos, and D. Jones-White. Undergraduate and graduate students' mental health during the covid-19 pandemic. 2020.

[12] T. Choudhury and A. Pentland. Characterizing social networks using the sociometer. *Proceedings of the North American association of computational social and organizational science (NAACSOS)*, 2004.

[13] V. Das Swain, H. Kwon, B. Saket, M. B. Morshed, K. Tran, D. Patel, Y. Tian, J. Philipose, Y. Cui, T. Plötz, et al. Leveraging wifi network logs to infer social interactions: A case study of academic performance and student behavior. *arXiv e-prints*, 2020.

[14] D. Dave, A. Friedson, K. Matsuzawa, J. J. Sabia, and S. Safford. Jue insight: Were urban cowboys enough to control covid-19? local shelter-in-place orders and coronavirus case growth. *Journal of urban economics*, page 103294, 2020.

[15] A. DePietro. Here's a look at the impact of coronavirus (covid-19) on colleges and universities in the u.s., 2020. Available at: https://www.forbes.com/sites/andrewdepietro/2020/04/30/impact-coronavirus-covid-19-colleges-universities/?sh=26c3f79c61a6.

[16] E. Dorn, B. Hancock, J. Sarakatsannis, and E. Viruleg. Covid-19 and student learning in the united states: The hurt could last a lifetime. *McKinsey & Company*, 2020.

[17] N. Eagle and A. S. Pentland. Reality mining: sensing complex social systems. *Personal and ubiquitous computing*, 10(4):255–268, 2006.

[18] M. H. S. Eldaw, M. Levene, and G. Roussos. Presence analytics: making sense of human social presence within a learning environment. In *2018 IEEE/ACM 5th International Conference on Big Data Computing Applications and Technologies (BDCAT)*, pages 174–183. IEEE, 2018.

[19] N. Ferguson, D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba Perez, G. Cuomo-Dannenburg, et al. Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand.

[20] N. M. Ferguson, D. A. Cummings, C. Fraser, J. C. Cajka, P. C. Cooley, and D. S. Burke. Strategies for mitigating an influenza pandemic. *Nature*, 442(7101):448–452, 2006.

[21] C. for Disease Control and Prevention. Community npis: Flu prevention in community settings, 2020. Available at: https://www.cdc.gov/nonpharmaceutical-interventions/community/index.html.

[22] C. for Disease Control and Prevention. People at increased risk and other people who need to take extra precautions, 2020. Available at: https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/index.html.

[23] S. Friedman, T. Hurley, T. Fishman, and P. Fritz. Covid-19 impact on higher education, 2020. Available at: https://www2.deloitte.com/us/en/pages/public-sector/articles/covid-19-impact-on-higher-education.html.

[24] R. Frutos, M. Lopez Roig, J. Serra-Cobo, and C. A. Devaux. Covid-19: the conjunction of events leading to the coronavirus pandemic and lessons to learn for future threats. *Frontiers in medicine*, 7:223, 2020.

[25] G. Gibson, J. S. Weitz, M. P. Shannon, B. Holton, A. Bryksin, B. Liu, S. Bramblett, J. Williamson, M. Farrell, A. Ortiz, C. T. Abdallah, and A. J. García. Surveillance-to-diagnostic testing program for asymptomatic sars-cov-2 infections on a large, urban campus - georgia institute of technology, fall 2020. *medRxiv*, 2021. Available at: https://www.medrxiv.org/content/early/2021/01/31/2021.01.28.21250700.

[26] P. T. Gressman and J. R. Peck. Simulating covid-19 in a university environment. *Mathematical biosciences*, 328:108436, 2020.

[27] E. S. Gurley. Strategies to support the covid-19 response in lmics, 2020. Available at: https://hopkinsglobalhealth.org/assets/documents/CGH_Webinar_-_Contact_Tracing_(Final_Version).pdf.

[28] M. Harris and K. Holley. Universities as anchor institutions: Economic and social potential for urban development. In *Higher education: Handbook of theory and research*, pages 393–439. Springer, 2016.

[29] N. Haug, L. Geyrhofer, A. Londei, E. Dervic, A. Desvars-Larrive, V. Loreto, B. Pinior, S. Thurner, and P. Klimek. Ranking the effectiveness of worldwide covid-19 government interventions. *Nature human behaviour*, 4(12):1303–1312, 2020.

[30] J. S. Jia, X. Lu, Y. Yuan, G. Xu, J. Jia, and N. A. Christakis. Population flow drives spatio-temporal distribution of covid-19 in china. *Nature*, 582(7812):389–394, 2020.

[31] R. E. Jordan, P. Adab, and K. Cheng. Covid-19: risk factors for severe disease and death, 2020.

[32] S. M. Kissler, C. Tedijanto, E. Goldstein, Y. H. Grad, and M. Lipsitch. Projecting the transmission dynamics of sars-cov-2 through the postpandemic period. *Science*, 368(6493):860–868, 2020.

[33] K. Kondo. Simulating the impacts of interregional mobility restriction on the spatial spread of covid-19 in japan. *medRxiv*, pages 2020–12, 2021.

[34] M. Korn. Colleges begin mapping out a more normal fall—with caveats, 2021. Available at: https://www.wsj.com/articles/colleges-begin-mapping-out-a-more-normal-fallwith-caveats-11615800600?reflink=desktopwebshare_permalink.

[35] W. H. Kruskal and W. A. Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

[36] M. E. Loades, E. Chatburn, N. Higson-Sweeney, S. Reynolds, R. Shafran, A. Brigden, C. Linney, M. N. McManus, C. Borwick, and E. Crawley. Rapid systematic review: the impact of social isolation and loneliness on the mental health of children and adolescents in the context of covid-19. *Journal of the American Academy of Child & Adolescent Psychiatry*, 2020.

[37] Z. Mehrab, A. goud Ranga, D. Sarkar, S. Venkatramanan, Y. C. Baek, S. Swarup, and M. Marathe. High resolution proximity statistics as early warning for us universities reopening during covid-19. *medRxiv*, 2020.

[38] S. S. Morse, R. L. Garwin, and P. J. Olsiewski. Next flu pandemic: what to do until the vaccine arrives? 2006.

[39] A. Nierenberg and A. Pasick. Schools briefing: University outbreaks and parental angst, 2020. Available at: https://www.nytimes.com/2020/08/19/us/colleges-closing-covid.html.

[40] G. D. of Public Health. Georgia department of public health daily status report, 2020. Available at: https://dph.georgia.gov/covid-19-daily-status-report.

[41] W. H. Organization et al. Critical preparedness, readiness and response actions for covid-19: interim guidance, 22 march 2020. Technical report, World Health Organization, 2020.

[42] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

[43] B. Pfefferbaum and C. S. North. Mental health and the covid-19 pandemic. *New England Journal of Medicine*, 383(6):510–512, 2020.

[44] R. C. Rabin. C.d.c. officials say most available evidence indicates schools can be safe if precautions are taken on campus and in the community, 2021. Available at: https://www.nytimes.com/2021/01/26/world/cdc-schools-reopening.html.

[45] A. Rodríguez, A. Tabassum, J. X. Jiaming Cui, J. Ho, P. Agarwal, B. Adhikari, and B. A. Prakash. Deepcovid: An operational deep learning-driven framework for explainable real-time covid-19 forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, Apr. 2021.

[46] R. Rojas and M. Delkic. As states reopen, governors balance existing risks with new ones, 2020. Available at: https://www.nytimes.com/2020/05/17/us/coronavirus-states-reopen.html.

[47] M. Salathé, M. Kazandjieva, J. W. Lee, P. Levis, M. W. Feldman, and J. H. Jones. A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*, 107(51):22020–22025, 2010.

[48] L. Y. Saltzman, T. C. Hansel, and P. S. Bordnick. Loneliness, isolation, and social support factors in post-covid-19 mental health. *Psychological Trauma: Theory, Research, Practice, and Policy*, 2020.

[49] P. Sapiezynski, A. Stopczynski, R. Gatej, and S. Lehmann. Tracking human mobility using wifi signals. *PloS one*, 10(7):e0130824, 2015.

[50] S. Saul and S. Hubler. Colleges vowed a safer spring. then students, and variants, arrived, 2021. Available at: https://www.nytimes.com/2021/02/09/us/colleges-covid.html.

[51] L. J. Schoen. Guidance for building operations during the covid-19 pandemic. *ASHRAE Journal*, 5(3), 2020.

[52] A. Smalley. Higher education responses to coronavirus (covid-19). In *National Conference of State Legislatures.[Accessed May 15, 2020]. https://www. ncsl. org/research/education/higher-education-responses-to-coronavirus-covid-19. aspx*, 2020.

[53] J. Suh, E. Horvitz, R. W. White, and T. Althoff. Population-scale study of human needs during the covid-19 pandemic: Analysis and implications. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*, pages 4–12, 2021.

[54] T. N. Y. Times. Tracking coronavirus cases at u.s. colleges and universities, 2021. Available at: https://www.nytimes.com/interactive/2021/us/college-covid-tracker.html.

[55] A. Trivedi, C. Zakaria, R. Balan, and P. Shenoy. Wifitrace: Network-based contact tracing for infectious diseasesusing passive wifi sensing. *arXiv preprint arXiv:2005.12045*, 2020.

[56] J. L. Wang and M. C. Loui. Privacy and ethical issues in location-based tracking systems. In *2009 IEEE International Symposium on Technology and Society*, pages 1–4. IEEE, 2009.

[57] S. Ware, C. Yue, R. Morillo, J. Lu, C. Shang, J. Kamath, A. Bamis, J. Bi, A. Russell, and B. Wang. Large-scale automatic depression screening using meta-data from wifi infrastructure. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(4):1–27, 2018.

[58] S. Watson, S. Hubler, D. Ivory, and R. Gebeloff. A new front in america's pandemic: College towns, 2020. Available at: https://www.nytimes.com/2020/09/06/us/colleges-coronavirus-students.html.

[59] K. A. Weeden and B. Cornwell. The small-world network of college classes: implications for epidemic spread on a university campus. *Sociological science*, 7:222–241, 2020.

[60] A. Wesolowski, C. Metcalf, N. Eagle, J. Kombich, B. T. Grenfell, O. N. Bjørnstad, J. Lessler, A. J. Tatem, and C. O. Buckee. Quantifying seasonal population fluxes driving rubella transmission dynamics using mobile phone data. *Proceedings of the National Academy of Sciences*, 112(35):11114–11119, 2015.

[61] C. Woolston. Signs of depression and anxiety soar among us graduate students during pandemic, 2020. Available at: https://www.nature.com/articles/d41586-020-02439-6.

[62] Y. Xue, I. S. Kristiansen, and B. F. de Blasio. Dynamic modelling of costs and health consequences of school closure during an influenza pandemic. *BMC public health*, 12(1):1–17, 2012.

[63] V. Yuen. Mounting peril for public higher education during the coronavirus pandemic. *Center for American Progress*, 2020.

[64] Y. Zhang, B. Li, and K. Ramayya. Learning individual behavior using sensor data: The case of gps traces and taxi drivers. *Forthcoming in Information Systems Research*, 2020.

[65] H. Zhao and Z. Feng. Staggered release policies for covid-19 control: Costs and benefits of relaxing restrictions by age and risk. *Mathematical biosciences*, 326:108405, 2020.

[66] Y. Zhou, R. Xu, D. Hu, Y. Yue, Q. Li, and J. Xia. Effects of human mobility restrictions on the spread of covid-19 in shenzhen, china: a modelling study using mobile phone data. *The Lancet Digital Health*, 2(8):e417–e424, 2020.

WiFi mobility models for COVID-19 enable less burdensome and more localized interventions for university campuses

Vedant Das Swain[*] Jiajia Xie[*,†] Maanit Madan[*]
Sonia Sargolzaei[*] James Cai[‡] Munmun De Choudhury[*]
Gregory D. Abowd[*,◇] Lauren N. Steimle[†] B. Aditya Prakash[*,1]

[*]College of Computing, Georgia Institute of Technology
[†]H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology
[‡]Department of Computer Science, Brown University
[◇]College of Engineering, Northeastern University

Supplementary Information

[1]To whom correspondence should be addressed: `badityap@cc.gatech.edu`

Supplementary Methods

In this section, first, we describe the primary data source for mobility models (WiMob), the data used for calibrating our simulations, and for comparison of contact networks with methods using enrollment data (En). Next, we describe how we construct counterfactual mobility networks under the two main policies of interest in our study: remote instruction (RI) and localized closures (LC). Finally, we describe an agent-based-model (ABM) of disease transmission, which has a contact structure based on WiMob, and how this model was calibrated.

Data

WiFi Mobility

We use data provided by the IT management facility at Georgia Institute of Technology (GT) which accumulates WiFi access point (AP) logs over time. The primary use of WiFi network logs is for maintenance and security purposes. We mine these logs post-hoc to describe the mobility of individuals on campus, which we refer to as WiMob. Here mobility is expressed by visits to certain locations that are demarcated by a corresponding AP. WiMob can also describe dwelling (duration of visits) and collocation (dwelling in the presence of others around the same AP).

The campus WiFi network spans 6959 APs distributed between 240 buildings (and some outdoor locations). We label APs according to which building they are inside, along with the closest room or space (e.g, hallway, lobby, suite, cafe, etc.). The AP may or may not reside inside the room, however, in most cases, only a single AP is associated with space. For less than 5% of the APs, the AP shared association to space with another AP. This many-to-one mapping is typically in the case of large halls and auditoriums. We resolve such many-to-one associations by using APs as a proxy of the space they are associated with. Therefore, individuals connected to different APs in the same space will still be identified as collocated. Similarly, an individual could connect to the network with multiple devices. However, less than 1% logs show that a user is connected to multiple APs around the same time. Therefore, WiMob is agnostic to which device connects to the APs to proxy the presence of the individual. For this study, we obtain the WiFi network logs retrospectively for all of Fall 2019, and the data for Fall 2020 was provided on a per-day basis. Each day, approximately $33,000$ different people connect their devices to the WiFi network on campus. Overall in Fall 2019, approximately $40,000$ different people connected to the campus network.

Note, that GT's managed WiFi network is not equipped with any Real-Time Location System (RTLS) [9, 27]. RTLS systems use Received Signal Strength Indicator (RSSI) values from multiple neighboring APs to provide high precise localization of individuals in terms of time and space. However, deploying such systems requires surveying the entire network. Additionally, precision localization raises more privacy concerns. These factors together

make it challenging for universities to justify the deployment of RTLS, unlike small retail settings that can monetize RTLS insights directly (e.g., insights on footfall can be tied to improving revenue).

Asymptomatic surveillance testing data

We calibrated the ABM using the publicly reported positivity rate on the GT campus as reported through the asymptomatic surveillance and diagnostic testing program [33]. The testing program used pooled saliva sample surveillance with follow-up diagnostic testing. The positivity rate was reported each day, but individuals must wait at least 1 week between tests. We aggregated the positivity rate by week during the Fall 2020 semester.

Confirmed case data

When calibrating our ABM, we considered the reported confirmed cases in Fulton County [40], the county in which GT is located. The 'Confirmed COVID-19 Cases' reported in this dataset are cases that have been confirmed with a positive molecular (PCR) test. We considered cases during the Fall 2020 semester to inform external transmissions in the ABM.

Enrollment network summary statistics

We compare structural properties of contact networks constructed with WiMob to contact networks constructed from GT's course enrollment transcripts (En) — the aggregate statistics as reported in [45]. The En network was based on Fall 2019 transcripts for GT's Atlanta campus. These were cleaned to account for cross-listed courses and was used to determine which students were classmates with each other to form a contact network.

WiFi Mobility Models

Inferring location from Logs

WiMob is our approach to describe contact between people and movement of people between locations. The first step requires using WiFi network logs to infer when individuals were at specific locations on campus by determining when devices were connected to the corresponding APs. Our system mines the WiFi network logs that are populated via the Simple Network Management Protocol (SNMP) — a standard and widely used monitoring protocol to organize device association behavior to a WiFi network. Periodic SNMP updates can be caused either by poll requests to the APs that log which devices are associated with it at that time. However, devices can appear invisible to detached from an AP for multiple reasons, for example, when devices are idle. Otherwise, SNMP updates can occur whenever a new device connects, which is typical when individuals move between APs. Our approach exploits this factor to first mine periods when individuals are moving, then identify periods of dwelling between movements, and finally determine collocation when two or more individuals are dwelling near the same AP. This system follows from other studies that mine WiFi

logs [18, 57] and the detailed processing pipeline and evaluation is presented in [13]. This system to infer collocations has been tested against lecture attendance and reports a high precision of 0.89, but a relatively lower specificity of 0.79 [13]. While it is not likely to show false-positives, it has a possibility to erroneously mark people absent from a location even though they were there. However, for the purposes of our study, a contact network is made over an entire day and it only needs a single collocation instance for us to consider contact. And therefore we believe this limitation would not significantly affect our models.

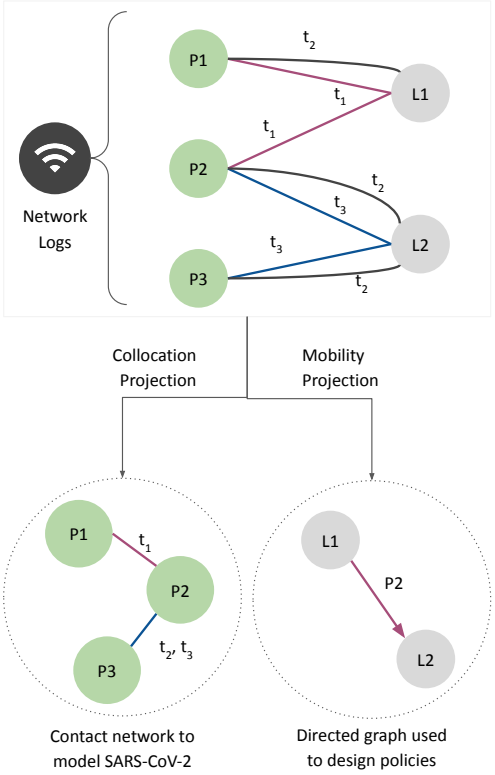

Figure S4: In a managed network, SNMP updates the logs by describing device association to an AP at a certain timestamp. WiMob mines these logs to characterize mobility as a bipartite graph. The nodes are partitioned to describe people nodes (e.g., $P1$, $P2$) connected to locations nodes (e.g., $L1$, $L2$). Every edge across the partition describes people visiting locations on campus during different times (e.g., t_1, t_2). Projecting the bipartite on people nodes helps construct a contact network (e.g., $P1$ and $P2$ were collocated at $L1$ at t_1), while projecting it on locations helps construct a directed movement graph ($P2$ dwelled at $L1$ and then at $L2$).

Characterizing Logs as Contact and Movement Networks

After inferring where an individual is located on campus, we represent the entire community behavior as graphs. We describe a bipartite graph, K, that shows when a user is at a given location on campus (Figure S4). This bipartite graph has edges connecting a set of m people, P, to a set of n locations, L. An individual can have multiple edges connecting to the location if they visited that location multiple times (e.g., t_1, t_2). The edge data contains the start

and end times of these dwelling periods. For these bipartite graphs, we make a projection on set P to describe collocation. This projection graph, G, contains an edge between users if they were visiting the same location during overlapping times. Since we do not use RTLS, our approach can only identify if people were in the vicinity of the same AP, but does not describe the distance between them. However, it can reasonably determine collocation in the same room [13]. Since our study is limited to localizing people indoors, we adapt the definition of *proximate contact* [27] where people might be "more than 6 feet but in the same room for an extended period". In our work, we use a lower bound threshold of 40 minutes to determine proximate contact. Therefore, individuals are only considered in contact when they are collocated in a room for 40 minutes or more. This threshold was set up to account for typical lecture duration on campus (for standard 3-credit hour courses taught 3 times a week). Every edge between two individuals contains a list of locations where they were possibly in contact. G forms the basis of the contact-network that we use an agent-based model to simulate. Alternatively, we also make a projection on the set L. This projection is a directed graph, H, where an edge from L_i to L_j represents movement from the first location to the next within a span of 60 minutes. GT's large urban campus with pedestrian pathways and motorized transit services enables direct movement between any two places on campus within the threshold. The 60 minutes threshold helps discount erroneously labeling returning from outside campus (e.g., non-residential students visiting two different locations between 2 days). H effectively describes how locations are connected and which locations could be more conducive to attracting and disseminating the virus. As a consequence, the H helps inform policy design. We compute the bipartite graph and its projections for each day of the semester.

Modeling Policies and Scenarios

RI: Offering Large Classes Online

As a response to COVID-19, prior work has recommended using EN to enforce a form of RI— moving classes large to an online remote instruction setup while other classes are offered in–person [26, 7, 59]. While we have access to aggregate insights on EN contact networks, our study protocol prohibits us from accessing course-specific information at an individual level. Therefore to infer individual enrollment, we analyze the edges of the bipartite graph K. For this, we first scrape the GT's course roster for Fall 2019 (filtered to only represent the Atlanta campus). This process provides us with a location and weekly schedule for every lecture conducted on campus, including its various sections. With this information, we are able to identify which edges represent visits to lectures, and subsequently, we can account for unique visitors to a lecture. Thus, we can first identify the number of unique individuals on campus who are enrolled in classes. The aggregate data from course enrollment reports that $21,299$ students were enrolled in Fall 2019. In comparison, our inference identifies $22,248$ students. The excess number can be explained by the fact that our method does

not distinguish between instructors, TAs, and students. Next, we study the unique visitors to every lecture in the scraped course schedule which gives us an estimate for the size of every class. Given the limitations of our data processing, actual enrollment sizes could be larger, but our process is less likely to count false positives [13]. Finally, to model RI, for the contact network G_t, we create a counterfactual network G'_t for each day t. These exclude collocations that took place at lecture locations during lecture times. If two people were connected solely by proximity during lectures — in a class with large enrollment — they will appear disconnected in the counterfactual network.

LC: Closing Important Locations

This article demonstrates the effectiveness of localized closures,LC, which are targeted interventions to seize mobility at different spaces on campus. For this, we identify important locations on campus by analyzing H. In the main paper, LC uses *PageRank* [42] as an illustrative algorithm to identify important location nodes. For robustness, we apply various additional algorithms to identify highly authoritative nodes in H — betweenness centrality [15], eigenvector centrality [5], and load centrality [29]. In the SI Appendix, we distinguish these different policies as $\mathsf{LC}_{\mathrm{PRank}}$, $\mathsf{LC}_{\mathrm{BCen}}$, $\mathsf{LC}_{\mathrm{ECen}}$, $\mathsf{LC}_{\mathrm{LCen}}$. Since RI captures a weekly schedule to determine enrollment, LC is implemented to find locations based on behavior from the past 7 days of mobility. We apply the weighted version of the algorithms mentioned earlier on the directed graph representing movement, H. The edge weight is based on the number of instances of movement between any L_i and L_j. After sorting the locations by importance, we determine the number of locations to shut down based on different budgets induced by RI— mobility and risk of exposure. For this purpose, we take the approach of a greedy algorithm which successively removes highly-ranked locations till the constraint is met (within 1% margin of error). Similar to RI, LC also render counterfactual collocation networks, $G"_t$ for each day t. In these networks, we remove instances of collocations that occurred at the shutdown locations. Figure S22 and Figure S23 shows the categories of buildings where different spaces are closed by LC policies.

Inducing Budgets and Characterizing Behavioral Scenarios

We now describe how we compare the RI and LC policies. First, we consider the effects of these policies under three behavioral scenarios. These scenarios express the spillover effects of closure that lead to students avoiding campus entirely because their entire schedule is forced online. This analysis assumes that the motivation to be present on campus is determined primarily by enrollment. We consider that, if a student has a full course load (enrolled in a minimum of 3 classes) and all their classes are offered online, that student might have less incentive to visit campus at all (for any engagement) and thus practice *Avoidance*. Since LC could end up closing classrooms, it can also lead to academic schedules being affected and elicit *Avoidance* behavior. As a result, we describe three scenarios. *Persistence*, is the preliminary, or null scenario, which represents no *Avoidance*. This counterfactual collocation graph only removes edges directly affected by RI or LC. The second scenario

we model is *Non-Residential Avoidance* where only non-residential students with full online schedules stop visiting campus entirely. Here the counterfactual graph will remove all edges of non-residential students with fully online schedules. Lastly, the third scenario we model is *Complete Avoidance* where any student with fully online schedules stops activity on campus entirely (including residential students). Here the counterfactual graph will remove all edges from any student with fully online schedules. Since our study protocol prohibits us from mapping our data to other sources, we heuristically infer which individuals are likely to be residential and which are not. We label individuals as residential when they dwell an average of at least 15 minutes at residential locations between 6pm and 10am, on workdays (Monday–Thursday).

Under each scenario, we limit the number of locations that can be closed under the LC policy to ensure the level of restriction is constrained to be similar to the RI policy. We limit the number of locations under two types of restrictive budgets. The first budget is based on *mobility*, which is the percentage of edges remaining in the bipartite graph if a policy were to be implemented. The second budget is based on *exposure risk*, which is the number of unique individuals who would be in the 1-hop collocation neighborhood of positive individuals. We compute this budget by randomly sampling 2.5% of the population as positive, based on the highest 7-day average positivity rate reported by GT [25] in Fall 2019. Note, however, the effect of RI on campus can vary in different behavioral scenarios, thereby changing the budget available to design a comparable LC policy. For instance, the number of people at exposure risk is much lower in *Complete Avoidance*. As a result, we build multiple alternate networks representing the effect of policies under counterfactual behavioral scenarios.

The infection reduction outcomes and burdens of different policy interventions (under various scenarios and budgets) is described in Table S5—Table S8 presents boxplots that compares the distribution of disease control outcomes. Figure S14—Figure S17 show cumulative plots of disease control outcomes.

Agent-Based Model & Simulation

Agent-Based Model

We constructed an agent-based model (ABM) that captures the spread of COVID-19 between individuals active within the GT community. The model is used to evaluate the effectiveness of different policy interventions. We consider a modified version of the SEIR framework for simulating the spread of COVID-19 [39, 10] by using an underlying contact network given by WiMob. Figure S5 shows the compartments of the framework. The *susceptible* state (S) represents individuals who have not been infected and can contract the disease by having contact with an infectious individual. The *exposed* state (E) is canonically equivalent to the "incubation period" and is similar to the pre-symptomatic state found in related work [44, 20]. Individuals are considered *infectious* when they are in either the *asymptomatic* state (Asym) or *symptomatic* state (Sym). Individuals in the *asymptomatic*

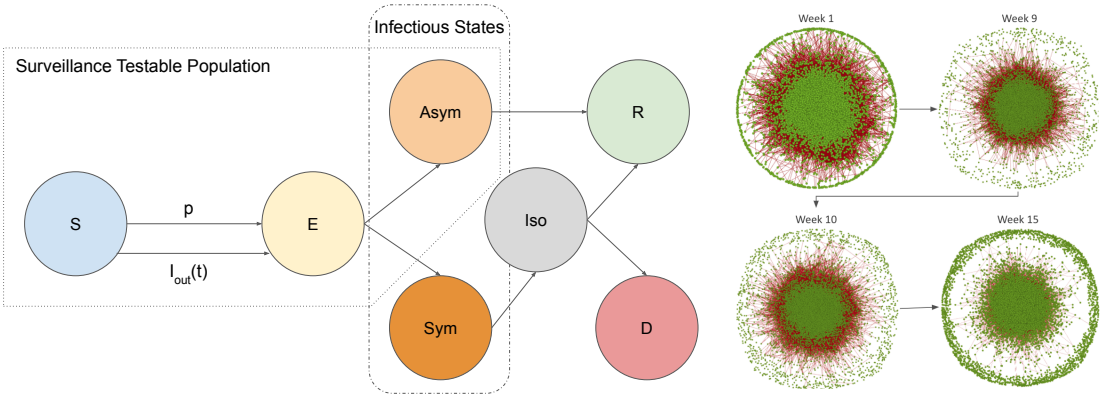

(a) (b)

Figure S5: (a) The schematic of the compartments in our modified SEIR model. By the design of the GT surveillance testing [33, 25], the total testable population is defined as the summation of *susceptible*, *exposed*, and *asymptomatic*. Infectious persons are in either *symptomatic* or *symptomatic*. For every effective edge in the mobility network, a susceptible individual that is exposed to an infectious person becomes infected with probability p. Individuals may also get infected due to an exposure not captured by the WiMob network which occurs with probability $I_{out}(t)$ on day t. account for new infected cases. (b) The mobility behavior represented by WiMob changes every day of the semester (shown weekly here). The contact network constructed from WiMob forms the underlying contact structure of the ABM.

state are assumed to be the major "spreaders" [20] and transmit the infections to *susceptible* individuals before they are *recovered* (R) [26] — after 7 days [20]. Since *asymptomatic* is considered a state of mild severity [37], individuals in this state do not have a risk of fatality. By contrast, for individuals in the *symptomatic* state, will be eventually *isolated* (Iso) (e.g. self-quarantine, or hospitalization on campus). Once in the *isolated* state, they cannot transmit the disease to individuals in the *susceptible* state. Unlike the *asymptomatic* track, the *symptomatic* state is considered critical severity. Therefore, after moving to the *isolated* state, individuals have risk of fatality and entering the death state (D). If the *isolated* individual survives, they enter the *recovered* state. We assume immunity is preserved and therefore after recovery the individual is no longer *susceptible*.

Definitions

Let $t = \{0, 1, 2, 3, ..., T\}$ be the index of days in simulations. We denote the sequence of dynamic collocation networks indexed by day t, as $\{G_t(A_t, B_t)\}_{t=0}^{T}$. A_t is the set of vertices, i.e. individuals on campus, and B_t is the set of edges. The universe set of the population throughout the simulation time period is given by $M = \bigcup_{i=1}^{T} A_t$. For convenience, we use $a_i \in M$ to index every person in the universe population set.

The SEIR model consists of seven compartments. Each of these corresponds to a function of population subsets with respect to day t: *susceptible* S(t), *exposed* E(t), *asymptomatic* *Asym*(t), *symptomatic* *Sym*(t), *isolation* I(t), *recovered* R(t), and *dead* D(t). For example, $a_i \in I(t)$ means a_i is in the *isolation* state at day t. We use $N^t_{S\to E}$, $N^t_{E\to Asym}$, $N^t_{E\to Sym}$,

$N^t_{Asym\to R}$, $N^t_{Sym\to I}$, $N^t_{I\to R}$, and $N^t_{I\to D}$ to denote the transitions between states between day t and day $t+1$.

Model Initialization

The entire population M is fixed where $M = S(t)+E(t)+Asym(t)+Sym(t)+I(t)+R(t)+D(t)$ for all t. To capture the positivity out of the students coming back to campus at the start of the semester, we initialize the system by setting a subset of M into $Asym(0)$ and the reminder into $S(0)$. The initial percentage of *asymptomatic* is described by:

$$Asym(0) \sim Binomial(M, I_0)$$

$$S(0) \sim M - Asym(0)$$

where I_0 is a parameter defined as the initial percentage of *Asymptomatic* at day $t = 0$.

New exposures

We consider two ways that an individual in the ABM could be exposed: (i) exposures that occur due to contacts among individuals captured by the mobility network (*internal transmission*) and (ii) exposures that occur due to contacts that occur outside of the mobility network (*external transmission*).

Internal transmissions happen exclusively among individuals in the model. On any given day, an edge becomes effective, when one of the *susceptible* individual comes in contact with the other which is infectious, i.e. *asymptomatic* or *symptomatic*, individual. Therefore, for every effective edge between two such people, the probability of the susceptible individual getting *exposed* is described by the transmission probability p, which is another model parameter. The probability for an susceptible individual a_i entering *exposed* at the end of day t is given by the following function:

$$\mathrm{f}_p(a_i, t, p) = \begin{cases} 1-(1-p)^{e(t,a_i)}, & \text{if } a_i \in V_t \\ 0, & \text{otherwise} \end{cases}$$

Here, $e(t, a_i)$ is the number of effective edges of individual a_i at time t. Since $(1-p)^{e(t,a_i)}$ is the probability that a_i does not contracted the disease at time t under $e(t, a_i)$ Bernoulli trials, $1-(1-p)^{e(t,a_i)}$ is the probability that at least one effective edge leading a_i to *exposed*.

In addition to exposure due to internal transmission, we also consider new exposure due to external transmission. We consider external transmission to be exposure resulting from the physical collocations outside the scope of mobility network. For instance, the WiMob does not capture the connections between individuals without access to the campus WiFi or

someone contacting infectious persons outside the campus. To reflect this risk in our model, for any day t, $I_{out}(t)$ describes the probability of infection on day t from a collocation that is external to the mobility network. We assume that the probability an individual is infected due to an external source is proportional to the number of cases in the broader community. Therefore, we model the probability of external infection as a function of confirmed cases in Fulton county, where GT is located [40]. C_t represents the confirmed cases reported by Fulton County where C_{max} is the maximum number of the cases over the whole period, $I_{out}(t)$ is given by

$$I_{out}(t) = \alpha * \frac{C_t}{C_{max}}$$

where α is a parameter scaling the normalized confirm cases in the surrounding county. The resulting number of external infections on day t is then modeled to be are Binomial with $|S(t)|$ trials with probability of success $I_{out}(t)$.

In summary, for every day $t > 0$, the overall number of individuals that become newly exposed is represented as $N^t_{S\to E}$ which is the result of both external and internal transmissions.

$$\mathrm{N}^t_{S\to E} \sim \underbrace{Binomial(|S(t)|, I_{out}(t))}_{\text{external infections}} + \underbrace{\sum_{a_i \in M} f_p(a_i, t, p)}_{\text{internal transmissions}}$$

Model dynamics after exposure

After exposure, individuals in the model will progress through other disease states in our model. We update the number of individuals in each state daily to reflect transitions between them. The transitions between the states on day t are summarized according to the following equations:

$$\mathrm{S(t+1)\ -\ S(t)} = -\ \mathrm{N}^t_{S\to E}$$

$$\mathrm{E(t+1)\ -\ E(t)} = \mathrm{N}^t_{S\to E} - N^t_{E\to \mathrm{Asym}} - N^t_{E\to \mathrm{Sym}}$$

$$Asym\mathrm{(t+1)\ -\ }Asym\mathrm{(t)} = \mathrm{N}^t_{E\to \mathrm{Asym}} - N^t_{\mathrm{Asym}\to R}$$

$$Sym\mathrm{(t+1)\ -\ }Sym\mathrm{(t)} = \mathrm{N}^t_{E\to \mathrm{Sym}} - N^t_{\mathrm{Sym}\to I}$$

$$\mathrm{I(t+1)\ -\ I(t)} = \mathrm{N}^t_{\mathrm{Sym}\to I} - N_{I\to D} - N_{I\to R}$$

R(t+1) - R(t) = $\mathrm{N}_{I\to R}$

D(t+1) - D(t) = $\mathrm{N}_{I\to D}$

After an individual has been exposed, they will spend Δ_{S} days in an incubation period. At day Δ_{S} after their exposure, individuals will become a *symptomatic* infection with probability p_S. Otherwise the agent will become an *asymptomatic* infection This process is given by the following two equations:

$$\mathrm{N}^t_{E\to\mathrm{Sym}} \sim \begin{cases} Binomial(|E(t-\Delta_{\mathrm{S}})|, p_S), & t \geq \Delta_{\mathrm{S}} \\ 0, & \text{otherwise} \end{cases}$$

$$\mathrm{N}^t_{E\to\mathrm{Asym}} \sim \begin{cases} |E(t-\Delta_{\mathrm{S}})| - N^t_{E\to\mathrm{Sym}}, & t \geq \Delta_{\mathrm{S}} \\ 0, & \text{otherwise} \end{cases}$$

Individuals who enter the *asymptomatic* state will recover after $\Delta_{\mathrm{Asym}\to R}$ days since they were first *exposed*. Thus, we represent the number of transitions from *asymptomatic* to *recovered* on day t as:

$$\mathrm{N}^t_{\mathrm{Asym}\to R} \sim \begin{cases} N^{t-\Delta_{\mathrm{Asym}\to R}}_{E\to\mathrm{Asym}}, & t \geq \Delta_{\mathrm{Asym}\to R} \\ 0, & \text{otherwise} \end{cases}$$

On the other hand, individuals who enter the *symptomatic* will eventually enter the *isolation* state [20]. The time that individuals spend in the *symptomatic* state before entering the *isolated* state is normally distributed $\delta^t_I \sim Normal(\Delta_I, \sigma^2_I)$. We simulate each individual's transition between *symptomatic* and isolated by using a sampling function $\Gamma(a_i, t, \Delta_t)$ and a function $\tau(a_i, t)$ that returns the days since exposed respectively:

$$\Gamma(a_i, t, \delta^t_I) = \begin{cases} 1, & t - \tau(a_i, t) \geq \delta^t_I \\ 0, & \text{otherwise} \end{cases}$$

$$\tau(a_i, t) = \begin{cases} \text{first day of } a_i \text{ entering } \textit{exposed}, & a_i \in Sym(t) \\ +\infty, & \text{otherwise} \end{cases}$$

Parameter	Definition	Value	Std	Source
p	Transmission probability: For any edge between a *susceptible* and *infectious* individual in the contact network, p is the probability that the *susceptible* person will enter into the *exposed* state. This only dictates internal transmission	0.034	0.007	Calibration
α	Scaling factor of the normalized confirmed cases in the surrounding county(()). This is the parameter for us to generate $I_{out}(t)$	0.032	0.0032	Calibration
I_0	Proportion of population that is *asymptomatic* at day 0	0.012	0.0009	Calibration
p_S	Probability of *exposed* persons becoming symptomatic	0.66	-	[20]
Δ_S	Incubation period (days) since the first day of exposure	5	-	[20]
$\Delta_{\text{Asym}\to R}$	Asymptomatic duration (days); it is the time taken for an *asymptomatic* person to recover since the first day of exposure	7	-	[20]
Δ_I, σ_I	Time of an *symptomatic* entering *isolated* since the first day of exposure of a *symptomatic* person	8	2	[16]
Δ_R, σ_R	Time for recovery for a *symptomatic*, since the first day of exposure	12	2	[22]
p_D	Death rate under isolation	0.0006	-	[22]

The variables p, α, and I_0 are estimated by calibrating the simulation model on the first 5 weeks of positivity rates provided by GT surveillance for Fall 2020, while incorporating external cases from Fulton County. These parameters were found by validating the ABM on the remaining weeks of Fall 2020. Figure S6 shows model estimate during the calibration and validation period.

Table S2: Model Parameters of the ABM

The aggregated transitions $N^t_{\text{Sym}\to I}$ between *symptomatic* and *isolated* is the sum of the distribution above on each day t.

$$N^t_{\text{Sym}\to I} \sim \sum_{a_i \in M} \Gamma(a_i, t, \delta^t_I)$$

Individuals who enter the *isolated* state may end up with one of two states: *dead* or *recovered*. We defined $N^t_{I\to D}$ as following another binomial distribution with parameter p_D:

$$N^t_{I\to D} \sim Binomial(|I(t)|, p_D)$$

The transitions between *isolation* and *recovered* is quite similar to the transitions between *symptomatic* and *isolation* except $\delta^t_R \sim Normal(\Delta_R, \sigma^2_R)$ where Δ_R and σ_R are the two parameters standing for the mean and standard deviation of days for an individual in the *isolation* state entering *recovered* since the first day of infection. This leads to:

$$N^t_{I\to R} \sim \sum_{a_i \in M} \Gamma(a_i, t, \delta^t_R).$$

Model calibration

Most of our model parameters can be estimated from previous studies (see Table S2). However, three parameters in our study are not easily estimated from previous studies: (1)

the proportion of the agents that begin the semester asymptotically infected, I_0, (2) the probability of transmission between a given infectious individual and susceptible individual given a contact in the mobility network, p, and (3) the scaling factor α used to determine probability of transmission due to contact outside of WiMob network on day t, $I_{out}(t)$ (see ()). We fit these three parameters to the published weekly positivity rate (percentage of asymptomatic cases) as reported by GT's asymptomatic surveillance testing program [33]. To fit the parameters, we performed calibration to minimize the root mean square of error(r.m.s.e) between the simulation estimates of the weekly positivity rate and the observed weekly positivity rate on GT's campus of the Fall 2020 semester as reported by the surveillance testing program.

To perform the calibration, we used two sets of public data pertaining to 2020 Fall semester at GT: (i) the confirmed cases in Fulton County [40], and (ii) the aggregated surveillance test positivity rate for each week [33]. The former helps estimate the daily external infection percentage. The latter is the ground truth trajectory we fit our model on. We consider the data aggregated by week because each individual on campus can only get tested once per week. The positivity rate provided by the surveillance testing data can be interpreted as the estimated percentage of new *asymptomatic* cases out of the total testable population which includes *susceptible*, *exposed*, and *asymptomatic* — with an assumption that every testable population get tested at the same rate.

To formalize the calibration problem, let R_w be the surveillance-testing aggregated result at week w. Let $S(I_0, \alpha, p, w)$ be the function of the simulation model which returns the percentage of new *asymptomatic* in week w out of the total testable population. For every combination of parameters, the predicted result for each week w is estimated by taking the average of N simulation outputs. The objective function is:

$$f(I_0, \alpha, p) = \sqrt{\frac{1}{W}\sum_{w=1}^{W}\left(\frac{\sum_{i=1}^{N} S(I_0\alpha, p, w)}{N} - R_w\right)^2}$$

The optimization problem is:

$$\min_{I_0,\alpha,p} f(I_0, \alpha, p)$$

We fit our model to the first 5 weeks of Fall 2020 and validate the results on the remaining weeks. After obtaining the optimal set of parameters, for robust comparison of policies with different viral variants, we generate a range of parameters by compromising the r.m.s.e within 40% of the minima [10]. First, we implement the Nelder Mead method [25] to discover the optimal set of parameters that minimizes the r.m.s.e. Next, we sample 40 different combinations of parameters within 40% of the minimum r.m.s.e to estimate the means and standard deviations of these parameters (Table S2). Throughout this paper, we pool together all simulation results across those parameters over multiple runs ($N = 15$) and report the 2.5th and 97.5th percentiles of the simulation outputs for every policy experiment.

Sensitivity Analyses

In this section, we design complementary experiments to inspect the robustness LC policies under different setups and calibration approaches. These variations are defined as follows:

- *Calibration periods (V1)*: For the results in the main paper, we discuss results with our ABM calibrated on the first 5 weeks of surveillance testing data. For additional analyses, the model parameters are re-estimated based on the surveillance data from week $5-9$ and $10-14$ in Fall 2020 at GT. The calibration is validated on the remaining weeks in the semester. Figure S6 shows the calibration and validation. The results of policy comparison with these variations can be found in Table S9 and Table S10, for weeks $5-9$ and $10-14$ respectively. Additionally, Figure S12 shows boxplots to compare the distributions of different policies, while Figure S18 and Figure S19 show cumulative plots of the disease control outcomes, for weeks $5-9$ and $10-14$ respectively.

- *Campuses and counties (V2)*: For the results in the main paper, the calibration of our ABM reflects certain latent factors inherent to GT that could affect both mobility behavior as well as testing results. To complement this we consider calibrating our data under different settings informed by surveillance testing from other similar large universities. This analysis is intended to represent the GT community in a different geographic setting, which is influenced by a different surrounding community, policies and resources. The new parameters are estimated based on the first 5 weeks of surveillance testing from the University of Illinois at Urbana-Champaign (UIUC) and the University of California, Berkeley (Berkeley) [31, 38], and the corresponding county data [12, 11] The calibration is validated on the remaining weeks in the semester. Figure S7 and Figure S8 show the calibration and validation for UIUC and Berkeley respectively. The results of policy comparison with these variations can be found in Table S11 and Table S12. Additionally, Figure S13 shows boxplots to compare the distributions of different policies, while Figure S20 and Figure S21 show cumulative plots of the disease control outcomes.

The estimated parameters with these calibration variations are described in Table S4. Both RI and LC are evaluated in the same infection reduction metrics and burden metrics again under scenarios S1, S2, and S3. Since the budgets are structural (mobility, and exposure risk) the LC policies are unchanged among the variants. Moreover, since the burden metrics are structural, those results are invariant.

Supplementary Discussion

Implications for Policy Design

To evaluate the efficacy of policies, we inspect infection reduction by simulating the disease with contact networks from Fall 2019. Since managed WiFi networks accumulate logs for

long periods of time, policymakers can use WiMob to model data from previous semesters and experiment with closure policies like LC. We show that WiMob can provide retrospective disease–mitigating insight into multiple counterfactual scenarios. For instance, policymakers can consider studying seasonal behaviors over multiple semesters for more robustness. Since the underlying data is longitudinal, it provides the flexibility to realistically assess policy interventions at different time points and also study updating policies. Restricting movement on campus at different time-points is known to exert varying degrees of control on disease spread [10]. Our data also shows that mobility on campus varies across the semester and therefore, allows policymakers to consider loosening shutdowns depending on the phase of the semester.

Policy design is determined by practical budgets. We model two kinds of budgets, mobility reduction and risk of exposure. The former represents disruptions in space utilization, availing services, and social life. The latter translates to the testing burden on campus. Our analysis determines the budget in different scenarios by observing the changes to the graph when large classes are moved online. This is to ensure an equitable comparison with targeted policies. However, in real situations these budgets can be relaxed or restricted based on that campus' preparedness to tackle a pandemic. For instance, a hypothetical campus that can test everyone every day might not be constrained by risk of exposure. Alternatively, policymakers can model other tangible budgets such as the capacity in isolation wards or available hospital beds. This can be informed by practical limitations of the campus. Similarly, this paper only assesses limited forms of cost, e.g., students avoiding campus or closing locations. From a financial perspective, university campuses can digitize their core service—education—but still realize losses from other curtailed services [23, 6, 58]. When students avoid campus it can lead to direct losses from meal passes and parking and also quantifiable losses to learning outcomes [2, 16] Policymakers can compute actual costs by complementing this data with information from other sources (e.g., revenue generated by cafes and stores on campus). This can help qualifying WiMob to reflect different costs and in turn help design policies that optimize for financial losses. Different campuses have different priorities and challenges in implementing policies.

Privacy, Ethics and Legal Considerations

We purposefully compare our prototype targeted policies against moving classes online because of practical budgets within the university. Both the WiMob and En based contact networks are derived from archival data accumulated by universities. This does not require instrumenting campus or its community with any new form of surveillance infrastructure. However, its use for a different purpose demands approval by an IRB. Moreover, acquiring these kinds of data would require collaborating with data-stewards (e.g., the IT department) to establish a data-use agreement. This document must clarify how the data will be de-identified, transferred, and stored.

For this form of data, the critical privacy challenge might not be localization itself, but rather the aggregation of data over a period of time [56]. Data spanning a longer period are more susceptible to cross-analyzing and identifying. To mitigate over-accumulation of data,

we suggest an adherence to principles of data minimization [36]. Instead of storing entire mobility graphs, the campus can compute and preserve only high-level insights, such as the importance of locations. This redacts any underlying individual behavior and corresponding identifiable information. Actually, for future purposes campuses can consider a form of differential privacy that authorizes limited forms of data querying depending on the privileges of the stakeholder [4].

An operational application would require the university to update the terms of use for its managed network. Particularly, the university should disclose how this data can be used in critical circumstances that invoke shared vulnerabilities [7]. On notifying the campus community of this change it offers individuals the choice to refrain from using the university network. Prior work on a sample within the same university campus shows that 90% of students are connected to the network on any given day [13]. Therefore, proposing such an opt-out condition can be viewed as an unfair choice. As a result, the campus needs to develop a contingency plan to accommodate network access to users who do not want their mobility behavior to constitute the aggregated insights.

Limitations and Future Work

This work presents evidence that university campuses can repurpose existing data sources to inform the design of LC policies that can control COVID-19. We evaluate these policies as alternatives to other data-driven, but, broad impact policies that universities consider implementing, such as moving large classes online. One of the drawbacks of this analysis, however, is that it assumes all edges to be the same. For example, when constraining by mobility, in real scenarios losing certain visits might be more valuable than others. Decline in mobility around profit-making services, such as shops and cafeterias, versus losing mobility at common rooms have a different tangible effects on campus. Currently we take an agnostic stance towards the mobility behavior, where all visits at all locations are the same. In reality, implementing policies could have inequitable qualitative impacts despite appearing to have similar network configuration. This can be improved by embedding more qualitative information into the network and conceiving ingenious ways to associate costs to edges.

Similar to the assumption that all visits and locations, the current work also assumes all people to be equal. However, different people have different underlying conditions that can make their vulnerabilities more concerning [43]. The privacy safeguards of this study restricted the research team from acquiring any additional demographic or historical information. Further work can attempt to characterize the nodes by randomly seeding the network to reflect the approximate demographic break up of the community. Alternatively, researchers could try to estimate some demographic based on behavior as well. However, to leverage accurate individual information, even for operational use during a public health emergency, policymakers and researchers need to develop new privacy protocols [28].

Lastly, this paper only studies three rudimentary scenarios, *persistence*, *non-residential avoidance*, *complete avoidance*. These scenarios assume that when a location is shutdown, the individuals who ought to have visited that location do not come into contact with anyone else during the same time. Yet, other substitution behaviors are possible and the richness of

networks leveraged with WiMob enables the exploration of various new scenarios that can be triggered by policy interventions on campus. For instance, individuals might not even visit transitory spaces, such as lobbies or cafes between classes. Certain collocations could be the consequence of social ties which might never be developed because of a shutdown (e.g., project teams meeting outside of class). Further research can illuminate the effects of policies in more specific scenarios by modeling post-intervention behavior more accurately.

	Cornell		Georgia Tech				
Contact Network	EN		EN		WIMOB		
Contact Situations	Course Lectures	RI	Course Lectures	RI	All Spaces	Course Lectures	RI
Number of Active Nodes	22051		21299		15379(±3353)	15379(±3353)	15380(±3353)
Average Contacts	529	22 − 41	341	30	152(±63)	86(±35)	86(±34)
Density	0.024	0.001	0.016	0.001	0.009(±0.002)	0.005(±0.001)	0.0053(±0.0014)
Largest Connected Component(%)	0.991	0.763	0.994	-	0.999(±0.001)	0.999(±0.02)	0.978(±0.025)
Average Shortest Path	2.47	3.75	2.54	3.54	2.67(±0.28)	3.26(±0.5)	2.953(±0.35)

We create a contact network of only students with WIMOB and compare it with insights from contact networks created with EN. On average, we find the contact network constructed with WIMOB shows fewer average contacts, lower density and higher average shortest path (between reachable paths). Moreover, within WIMOB itself, characterizing all spaces reveals more contacts and shorter paths than only focusing on contacts in lectures. While the proportion of the largest component appears similar, note that with WIMOB, on average about only 70% of the students visit campus on a given week. We further inspect the disease–mitigating structural changes of the RI policy on the network. We observe that the changes across all metrics with EN appear to be more drastic than compared to WIMOB.

Table S3: Comparison of Contact Network Structure (Fall 2019)

	Calibrating on Positivity Rate at GT			Calibrating on Positivity Rate with other University Behavior	
Parameter	weeks $0-4$	weeks $5-9$	weeks $10-14$	UIUC	Berkeley
p	0.034 ± 0.007	0.073 ± 0.005	0.0024 ± 0.0003	0.024 ± 0.0009	0.041 ± 0.003
α	0.032 ± 0.0032	0.0042 ± 0.0006	0.0159 ± 0.002	0.0069 ± 0.0013	0.038 ± 0.006
I_0	0.012 ± 0.0009	0.00057 ± 0.00007	0.0030 ± 0.0007	0.0039 ± 0.0013	0.0048 ± 0.0003
Optimal r.m.s.e	0.0034	0.0007	0.0015	0.0028	0.0031
Effective R_0 (min - max), Fall 2020	$1.15-1.18$	$1.17-2.14$	$0.33-0.95$	$1.12-1.19$	$1.24-1.28$
Effective R_0 (min - max), Fall 2019	$2.87-5.68$	$5.15-12.93$	$1.27-1.36$	$3.35-5.35$	$3.32-7.00$

The results in the main paper use variables p, α, and I_0 as estimated by calibrating the simulation model on the first 5 weeks of positivity rates provided by GT surveillance for Fall 2020, while incorporating external cases from Fulton County. For sensitivity analyses, we perform calibrations on GT data for weeks $5-9$ and $10-14$. Additionally, we perform calibrations on first five weeks of UIUC and Berkeley positivity rate (along with data from their respective county). These parameters were found by validating the ABM on the remaining weeks of Fall 2020. To assess the basic reproductive number (R_0) of our ABM we study the first 4 weeks of the disease. We find the effective R_0 to be higher for Fall 2019 than Fall 2020 as the mobility behaviors between the 2 semesters was vastly different. Note, Fall 2020 exhibits only 39% of the mobility we observe in Fall 2019. In fact, the ABM is calibrated on Fall 2020, where behavior was subject to pandemic related closures, but in Fall 2019 the mobility was not hindered by any interventions. Thus, Fall 2019 reflects a counterfactual of Fall 2020 without any closures.

Table S4: Calibration outcomes with variations

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	RI	LC		RI	LC		RI	LC	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	25.34(±12)	36.92(±14)**	34.30(±13)**	35.44(±10)	49.33(±11)**	52.19(±10)**	61.62(±7)	69.34(±5)**	64.44(±6)**
Total Infections (%)	6.99(±5)	10.63(±6)**	8.19(±5)**	14.88(±4)	13.96(±6)*	15.67(±6)	33.00(±5)	33.4(±5)	26.94(±5)**
Internal Transmissions (%)	17.13(±9)	22.62(±11)**	21.01(±11)**	27.58(±8)	35.35(±12)**	39.20(±11)**	54.00(±8)	70.89(±7)**	60.90(±9)**
Burdens on Campus									
Locations Affected	58	18	19	58	38	50	58	192	124
Students Avoiding (%)	0	0	0	9.30	0.20	0.45	27.21	12.45	6.57
Completely Isolated on Campus (%)	5.42	8.40	8.40	5.95	5.72	5.71	7.09	5.18	5.23

Note that this table is the same as Table 1. We repeat the results here for easier comparison of $\mathsf{LC}_{\mathrm{PRank}}$ to other algorithms shown in Table S6, Table S7 and Table S8. Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{PRank}}$ with RI. We find that $\mathsf{LC}_{\mathrm{PRank}}$ leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes are comparable in general but can vary by specific scenarios. In addition, every policy also exerts some burden on campus, either in terms of locations affected, students avoiding campus or isolation. We observe that $\mathsf{LC}_{\mathrm{PRank}}$ policies focus on fewer locations (except in *S3*). Moreover, these policies affect fewer student's schedules and therefore fewer people avoid campus due to completely remote schedules. Finally, $\mathsf{LC}_{\mathrm{PRank}}$ does not increase the percentage of people completely isolated on campus (p-value: < 0.01:*, < 0.001:**).

Table S5: Comparison of different $\mathsf{LC}_{\mathrm{PRank}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	RI	$\mathsf{LC}_{\mathrm{BCen}}$		RI	$\mathsf{LC}_{\mathrm{BCen}}$		RI	$\mathsf{LC}_{\mathrm{BCen}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	25.34(±12)	19.14(±12)**	30.93(±13)**	35.44(±10)	30.79(±13)**	51.87(±10)**	61.62(±7)	65.07(±6)**	61.38(±7)
Total Infections (%)	6.99(±5)	4.85(±4)**	7.74(±5)	14.88(±4)	7.76(±5)**	15.30(±6)	33.00(±5)	25.32(±5)**	22.08(±6)**
Internal Transmissions (%)	17.13(±9)	11.96(±9)**	19.64(±10)**	27.58(±8)	19.63(±10)**	38.74(±11)**	54.00(±8)	63.29(±8)**	54.00(±8)
Burdens on Campus									
Locations Affected	58	18	19	58	38	50	58	192	124
Students Avoiding (%)	0	0	0	9.30	0.07	0.45	27.21	11.47	6.74
Completely Isolated on Campus (%)	5.42	8.63	8.63	5.95	5.49	5.47	7.09	5.15	5.19

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{BCen}}$ with RI. We find that $\mathsf{LC}_{\mathrm{BCen}}$ leads to significantly improved peak infection reduction and internal transmission, when designed with the exposure risk budget, but can be worse with the mobility budget. In terms of reduction in total infections, the outcomes are typically worse. In addition, every policy also exerts some burden on campus, either in terms of locations affected, students avoiding campus or isolation. We observe that $\mathsf{LC}_{\mathrm{BCen}}$ policies focus on fewer locations (except in *S3*). Moreover, these policies affect fewer student's schedules and therefore fewer people avoid campus due to completely remote schedules. Finally, $\mathsf{LC}_{\mathrm{LCen}}$ does not increase the percentage of people completely isolated on campus (p-value: < 0.01:*, < 0.001:**).

Table S6: Comparison of different $\mathsf{LC}_{\mathrm{BCen}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	RI	$\mathsf{LC}_{\mathrm{ECen}}$		RI	$\mathsf{LC}_{\mathrm{ECen}}$		RI	$\mathsf{LC}_{\mathrm{ECen}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	25.34(±12)	36.15(±13)**	36.13(±13)**	35.44(±10)	44.52(±12)**	51.33(±10)**	61.62(±7)	65.13(±6)**	62.15(±7)
Total Infections (%)	6.99(±5)	8.66(±6)**	8.69(±6)**	14.88(±4)	11.75(±6)**	14.96(±6)	33.00(±5)	25.39(±5)**	22.82(±6)**
Internal Transmissions (%)	17.13(±9)	22.33(±11)**	22.37(±11)**	27.58(±8)	29.95(±12)*	37.94(±11)**	54.00(±8)	63.56(±8)**	57.07(±10)**
Burdens on Campus									
Locations Affected	58	18	19	58	38	50	58	192	124
Students Avoiding (%)	0	0	0	9.30	0.20	0.55	27.21	13.11	6.96
Completely Isolated on Campus (%)	5.42	8.59	8.59	5.95	5.53	5.51	7.09	5.17	5.23

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{ECen}}$ with RI. We find that $\mathsf{LC}_{\mathrm{ECen}}$ leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes vary by specific scenarios. In addition, every policy also exerts some burden on campus, either in terms of locations affected, students avoiding campus or isolation. We observe that $\mathsf{LC}_{\mathrm{ECen}}$ policies focus on fewer locations (except in *S3*). Moreover, these policies affect fewer student's schedules and therefore fewer people avoid campus due to completely remote schedules. Finally, $\mathsf{LC}_{\mathrm{ECen}}$ does not increase the percentage of people completely isolated on campus (p-value: < 0.01:*, < 0.001:**).

Table S7: Comparison of different $\mathsf{LC}_{\mathrm{ECen}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	RI	$\mathsf{LC}_{\mathrm{LCen}}$		RI	$\mathsf{LC}_{\mathrm{LCen}}$		RI	$\mathsf{LC}_{\mathrm{LCen}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	25.34(±12)	22.42(±13)**	30.73(±13)**	35.44(±10)	32.85(±13)*	51.44(±10)**	61.62(±7)	65.01(±6)**	61.40(±7)
Total Infections (%)	6.99(±5)	5.48(±5)**	7.64(±5)	14.88(±4)	8.23(±5)**	15.03(±6)	33.00(±5)	25.33(±5)**	21.98(±6)**
Internal Transmissions (%)	17.13(±9)	13.79(±9)**	19.37(±10)**	27.58(±8)	20.86(±11)**	38.08(±11)**	54.00(±8)	63.28(±8)**	55.28(±9)
Burdens on Campus									
Locations Affected	58	18	19	58	38	50	58	192	124
Students Avoiding (%)	0	0	0	9.30	0.07	0.43	27.21	11.47	6.73
Completely Isolated on Campus (%)	5.42	8.63	8.63	5.95	5.49	5.47	7.09	5.15	5.20

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{LCen}}$ with RI. We find that $\mathsf{LC}_{\mathrm{LCen}}$ leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes are comparable in some scenarios but can vary in specific scenarios. In addition, every policy also exerts some burden on campus, either in terms of locations affected, students avoiding campus or isolation. We observe that $\mathsf{LC}_{\mathrm{LCen}}$ policies focus on fewer locations (except in *S3*). Moreover, these policies affect fewer student's schedules and therefore fewer people avoid campus due to completely remote schedules. Finally, $\mathsf{LC}_{\mathrm{LCen}}$ does not increase the percentage of people completely isolated on campus (p-value: < 0.01:*, < 0.001:**).

Table S8: Comparison of different $\mathsf{LC}_{\mathrm{LCen}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	Broad	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	20.10(±4)	25.60(±3)**	25.63(±3)**	31.25(±3)	42.32(±4)**	47.29(±4)**	62.35(±2)	88.87(±2)**	76.89(±3)**
Total Infections (%)	8.89(±2)	10.50(±3)**	9.70(±3)**	20.26(±2)	20.02(±3)	23.71(±4)**	46.72(±2)	67.92(±4)**	51.30(±4)**
Internal Transmissions (%)	9.97(±2)	11.51(±2)**	10.95(±2)**	21.84(±2)	22.51(±3)	26.64(±3)**	49.80(±2)	74.96(±3)**	56.89(±4)**

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{PRank}}$ with RI. We find that $\mathsf{LC}_{\mathrm{PRank}}$ leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes are better in general but can be comparable in specific scenarios. The burden exerted on campus is the same as structural impacts of $\mathsf{LC}_{\mathrm{PRank}}$ (Table S5). (p-value: < 0.01:*, < 0.001:**).

Table S9: Comparison of different $\mathsf{LC}_{\mathrm{PRank}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $5-9$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	Broad	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	−1.75(±8)	3.65(±8)**	−1.95(±8)	3.88(±8)	−2.24(±8)**	−2.06(±8)**	20.39(±7)	7.57(±8)**	2.81(±8)**
Total Infections (%)	3.93(±9)	10.36(±8)**	5.13(±9)	9.87(±8)	6.36(±9)**	6.48(±9)**	26.02(±7)	16.37(±8)**	11.80(±8)**
Internal Transmissions (%)	42.33(±10)	61.15(±7)**	56.25(±8)**	49.83(±9)	67.10(±6)**	69.10(±6)**	74.74(±5)	84.80(±3)**	79.90(±4)**

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{PRank}}$ with RI. We find that $\mathsf{LC}_{\mathrm{PRank}}$ leads to significantly improved peak infection reduction and internal transmission. In terms of reduction in total infections, the outcomes are better in general but can be comparable in specific scenarios. The burden exerted on campus is the same as structural impacts of $\mathsf{LC}_{\mathrm{PRank}}$ (Table S5). (p-value: < 0.01:*, < 0.001:**).

Table S10: Comparison of different $\mathsf{LC}_{\mathrm{PRank}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $10-14$ in Fall 2020 at GT

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	Broad	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	41.40(±3)	60.44(±2)**	59.52(±2)**	49.75(±2)	74.22(±2)**	76.44(±2)**	78.14(±1)	85.81(±1)**	83.71(±1)**
Total Infections (%)	18.46(±3)	27.12(±3)**	25.25(±3)**	27.09(±3)	38.00(±4)**	40.68(±4)**	51.97(±3)	59.93(±5)**	54.07(±5)**
Internal Transmissions (%)	28.22(±3)	40.93(±3)**	39.09(±3)**	37.89(±3)	58.47(±2)**	65.45(±2)**	68.04(±2)	86.45(±1)**	80.08(±1)**

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{PRank}}$ with RI. We find that $\mathsf{LC}_{\mathrm{PRank}}$ leads to significantly improved peak infection reduction, internal transmission and total infections. The burden exerted on campus is the same as structural impacts of $\mathsf{LC}_{\mathrm{PRank}}$ (Table S5). (p-value: < 0.01:*, < 0.001:**).

Table S11: Comparison of different $\mathsf{LC}_{\mathrm{PRank}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at UIUC

Scenario	S1: Persistence			S2: Non-Res Avoidance			S3: Complete Avoidance		
Policy	Broad	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$		RI	$\mathsf{LC}_{\mathrm{PRank}}$	
Budget	-	Mobility (95.5%)	Exposure Risk (18800)	-	Mobility (92.3%)	Exposure Risk (16900)	-	Mobility (69.2%)	Exposure Risk (12700)
Infection Reduction Outcomes									
Peak Infections (%)	29.13(±3)	36.46(±5)**	36.34(±5)**	38.83(±3)	54.95(±4)**	58.88(±4)**	66.69(±2)	78.18(±1)**	77.65(±2)**
Total Infections (%)	6.34(±3)	8.59(±3)**	7.28(±3)**	14.71(±3)	13.18(±4)**	14.83(±4)	33.86(±4)	33.98(±5)	27.10(±5)**
Internal Transmissions (%)	15.99(±3)	20.43(±4)**	19.17(±4)**	27.01(±3)	34.60(±4)**	38.78(±4)**	55.01(±2)	74.65(±2)**	63.57(±3)**

Within each scenario, we perform the Kruskal-Wallis H-Test [35] to compare outcomes of $\mathsf{LC}_{\mathrm{PRank}}$ with RI. We find that $\mathsf{LC}_{\mathrm{PRank}}$ leads to significantly improved peak infection reduction, internal transmission and total infections. The burden exerted on campus is the same as structural impacts of $\mathsf{LC}_{\mathrm{PRank}}$ (Table S5). (p-value: < 0.01:*, < 0.001:**).

Table S12: Comparison of different $\mathsf{LC}_{\mathrm{PRank}}$ policies in terms of controlling the disease and impacts on campus in Fall 2019; calibrated from week $0-4$ in Fall 2020 at UC Berkeley

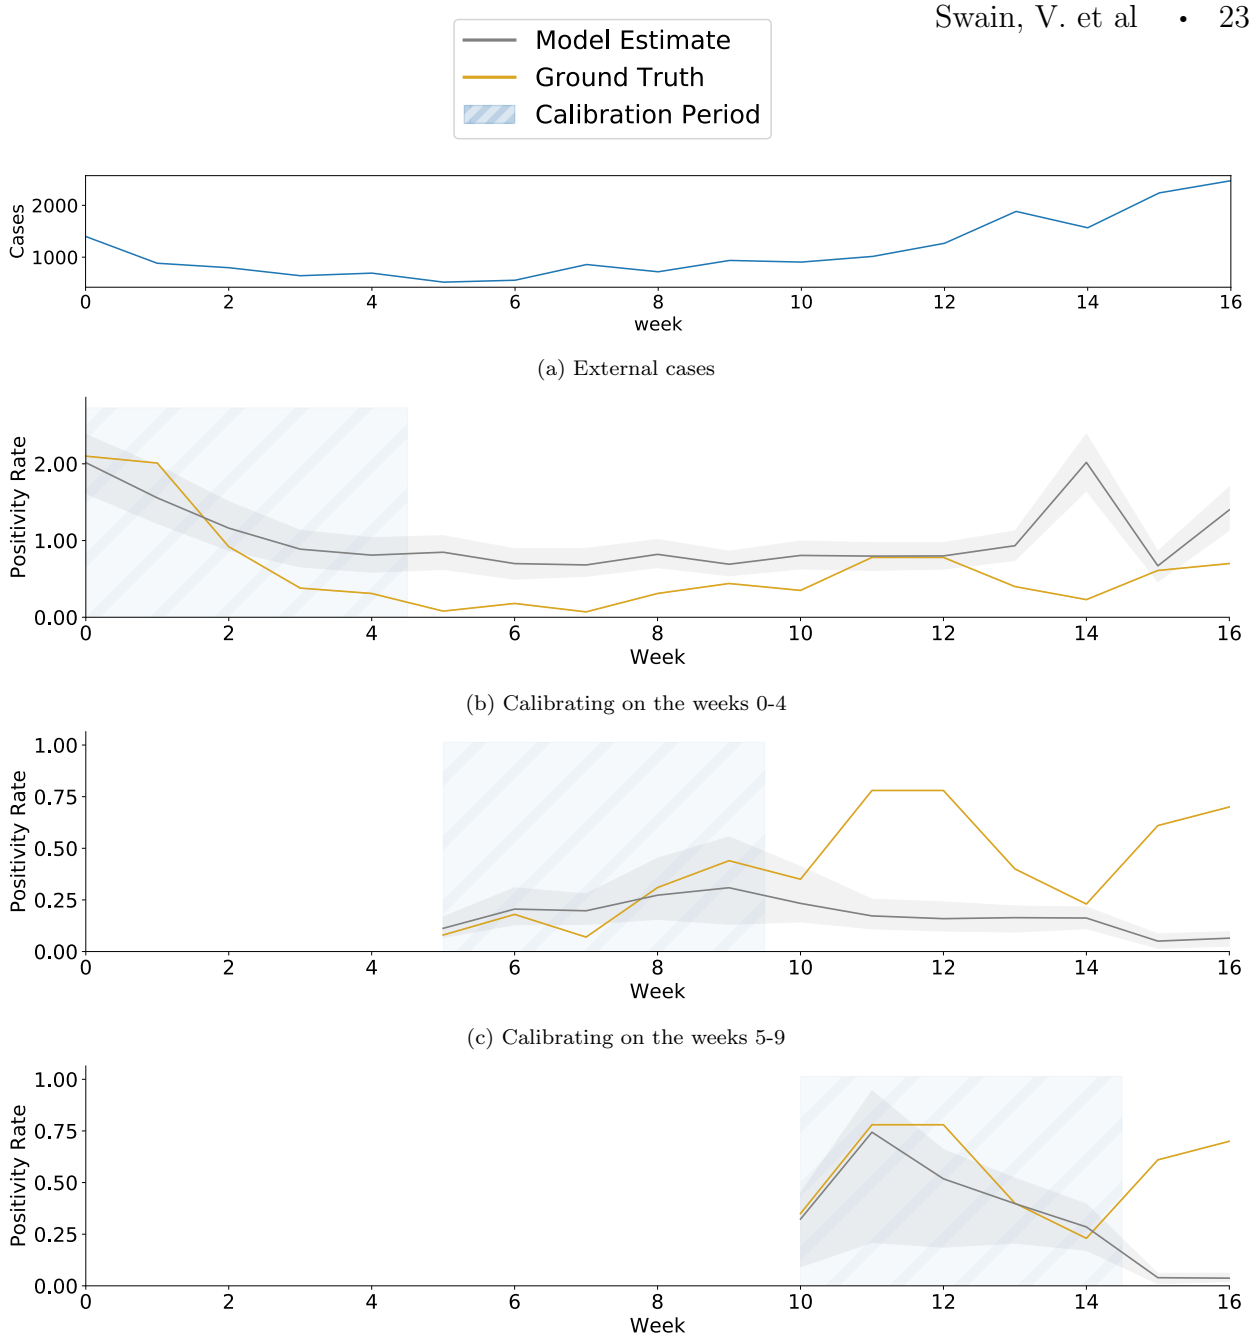

(a) External cases

(b) Calibrating on the weeks 0-4

(c) Calibrating on the weeks 5-9

(d) Calibrating on the weeks 10-14

Figure S6: We calibrate ABM on positivity rates from Fall 2020 at GT. The objective function of the calibration is to minimize the r.m.s.e. with the weekly average of positivity rate obtained from surveillance testing results at GT [25]. (a) The parameter that determines external transmission of infections on a given day, $I_{out}(t)$, is a function of cases in Fulton county (where GT is located). (b) The models discussed in the main paper are calibrated using the first 5 weeks of data. We illustrate the output for a range of parameters that incorporate quantitative uncertainty, i.e., within 40% of the r.m.s.e. (c, d) illustrate calibration on the second period of 5 weeks and third period of 5 weeks respectively. These only show the optimal parameter output. The shaded region around the lines show the 2.5^{th} and 97.5^{th} percentile.

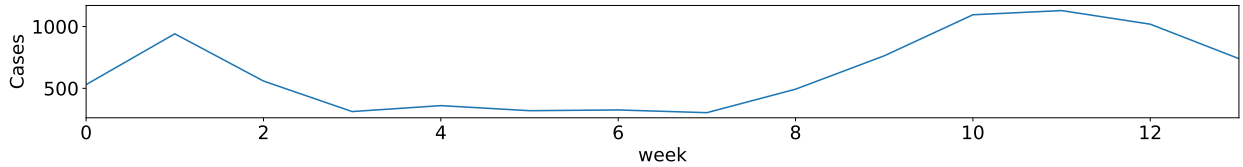

(a) External cases

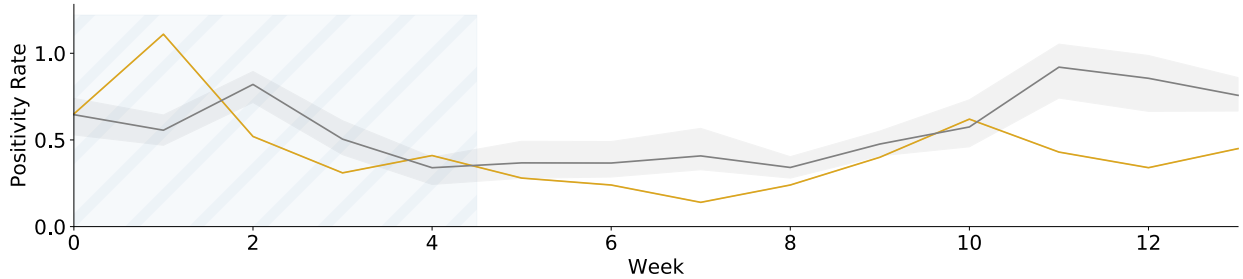

(b) Calibrating on the weeks 0-4

Figure S7: We calibrate ABM on positivity rates from first 5 weeks of Fall 2020 at UIUC. The objective function of the calibration is to minimize the r.m.s.e. with the weekly average of positivity rate obtained from surveillance testing results at GT [25]. (a) The parameter that determines external transmission of infections on a given day, $I_{out}(t)$, is a function of cases in Champaign county (where UIUC is located). (b) We illustrate the output for a range of parameters that incorporate quantitative uncertainty, i.e., within 40% of the r.m.s.e. The shaded region around the lines show the 2.5^{th} and 97.5^{th} percentile.

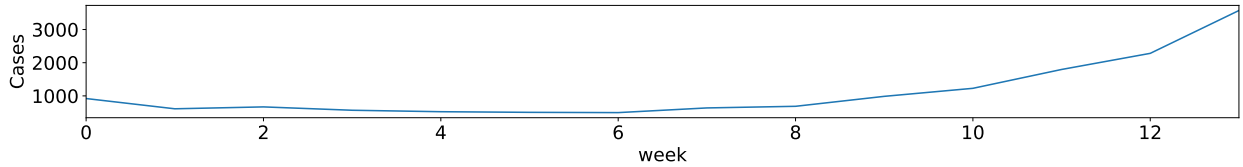

(a) External cases

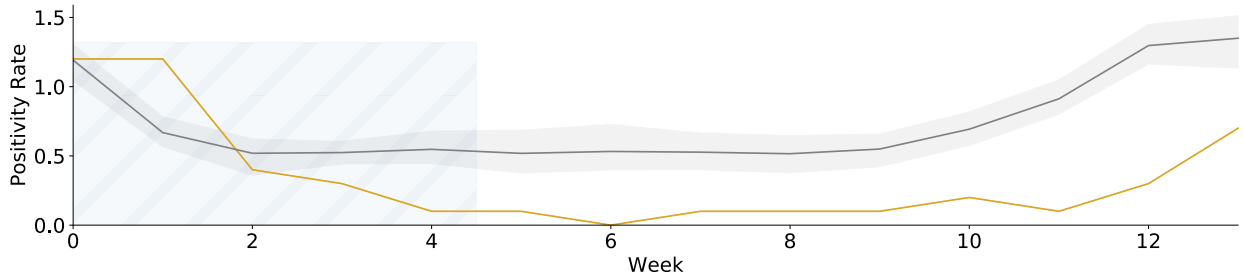

(b) Calibrating on the weeks 0-4

Figure S9: We calibrate ABM on positivity rates from first 5 weeks of Fall 2020 at UC Berkeley. The objective function of the calibration is to minimize the r.m.s.e. with the weekly average of positivity rate obtained from surveillance testing results at GT [25]. (a) The parameter that determines external transmission of infections on a given day, $I_{out}(t)$, is a function of cases in Alameda county (where UIUC is located). (b) We illustrate the output for a range of parameters that incorporate quantitative uncertainty, i.e., within 40% of the r.m.s.e. The shaded region around the lines show the 2.5^{th} and 97.5^{th} percentile.

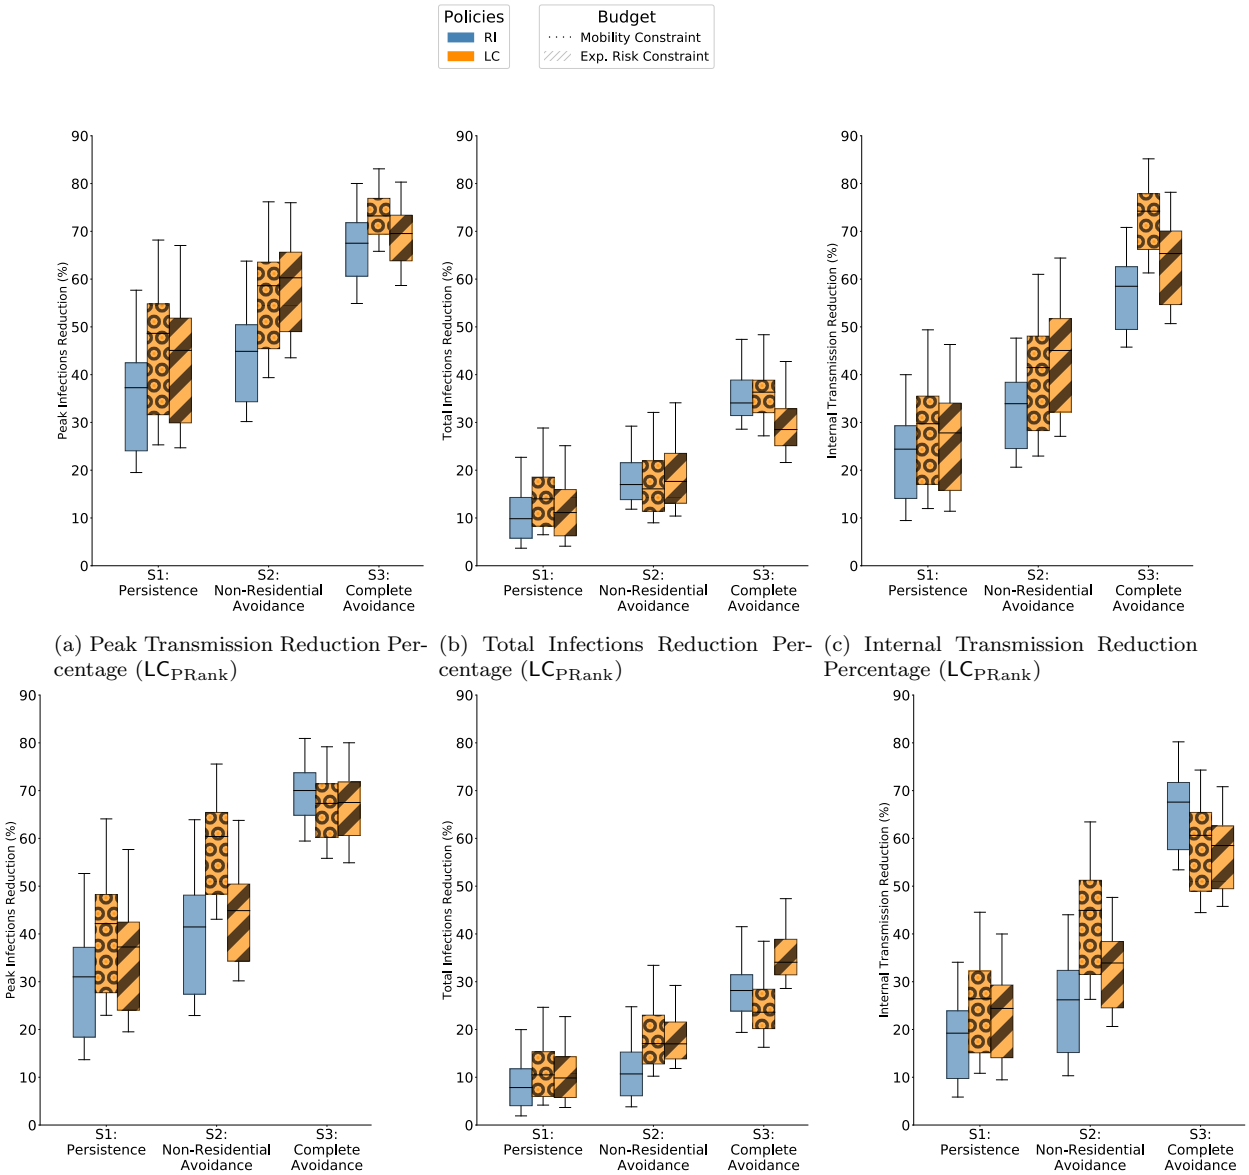

(a) Peak Transmission Reduction Percentage ($\mathsf{LC}_{\mathrm{PRank}}$)

(b) Total Infections Reduction Percentage ($\mathsf{LC}_{\mathrm{PRank}}$)

(c) Internal Transmission Reduction Percentage ($\mathsf{LC}_{\mathrm{PRank}}$)

(d) Peak Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{BCen}}$)

(e) Total Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{BCen}}$)

(f) Internal Transmission Reduction Percentage ($\mathsf{LC}_{\mathrm{BCen}}$)

Figure S10: Disease control outcomes in Fall 2019 for different algorithms of LC with the ABM is calibrated on weeks $0-4$ of Fall 2020 at GT. $(a-c)$ Comparison of RI with $\mathsf{LC}_{\mathrm{PRank}}$. Under all scenarios, for peak infection reduction (b) and internal transmission reduction (c), $\mathsf{LC}_{\mathrm{PRank}}$ shows better disease control outcomes than RI. For total infection reduction (b), $\mathsf{LC}_{\mathrm{PRank}}$ is better in *S1*, worse in *S3* when designed within an exposure risk budget, and comparable in others. $(d-f)$ Comparison of RI with $\mathsf{LC}_{\mathrm{BCen}}$. Under all scenarios, for peak infection reduction (d) and internal transmission reduction (f) $\mathsf{LC}_{\mathrm{BCen}}$ is better when designed within an exposure risk budget. For total infection reduction (e), $\mathsf{LC}_{\mathrm{BCen}}$ is always worse than RI

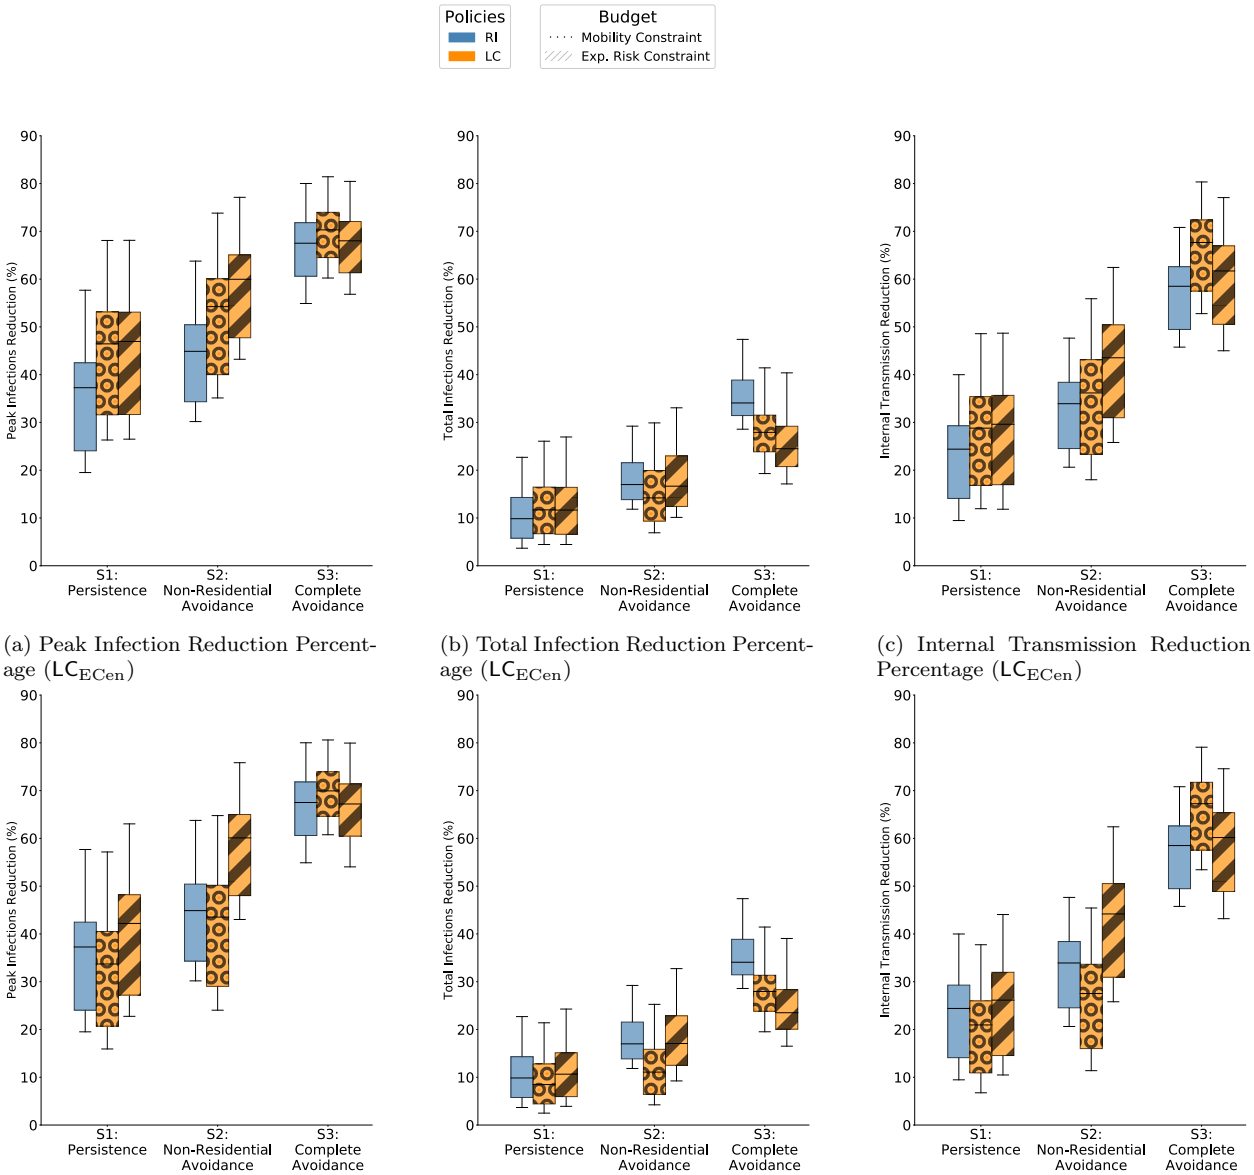

(a) Peak Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{ECen}}$)

(b) Total Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{ECen}}$)

(c) Internal Transmission Reduction Percentage ($\mathsf{LC}_{\mathrm{ECen}}$)

(d) Peak Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{LCen}}$)

(e) Total Infection Reduction Percentage ($\mathsf{LC}_{\mathrm{LCen}}$)

(f) Internal Transmission Reduction Percentage ($\mathsf{LC}_{\mathrm{LCen}}$)

Figure S11: Disease control outcomes in Fall 2019 for different algorithms of LC with the ABM is calibrated on weeks $0-4$ of Fall 2020 at GT. $(a-c)$ Comparison of RI with $\mathsf{LC}_{\mathrm{ECen}}$. Under all scenarios, for peak infection reduction (b) and internal transmission reduction (c), $\mathsf{LC}_{\mathrm{ECen}}$ shows better disease control outcomes than RI. For total infection reduction (b), $\mathsf{LC}_{\mathrm{ECen}}$ is better in *S1* and worse in *S3* when designed within an exposure risk budget. $(d-f)$ Comparison of RI with $\mathsf{LC}_{\mathrm{ECen}}$. Under all scenarios, for peak infection reduction (d) and internal transmission reduction (f), $\mathsf{LC}_{\mathrm{ECen}}$ shows better disease control outcomes than RI. For total infection reduction (e), $\mathsf{LC}_{\mathrm{ECen}}$ is better in *S1* and worse in *S3* when designed within an exposure risk budget.

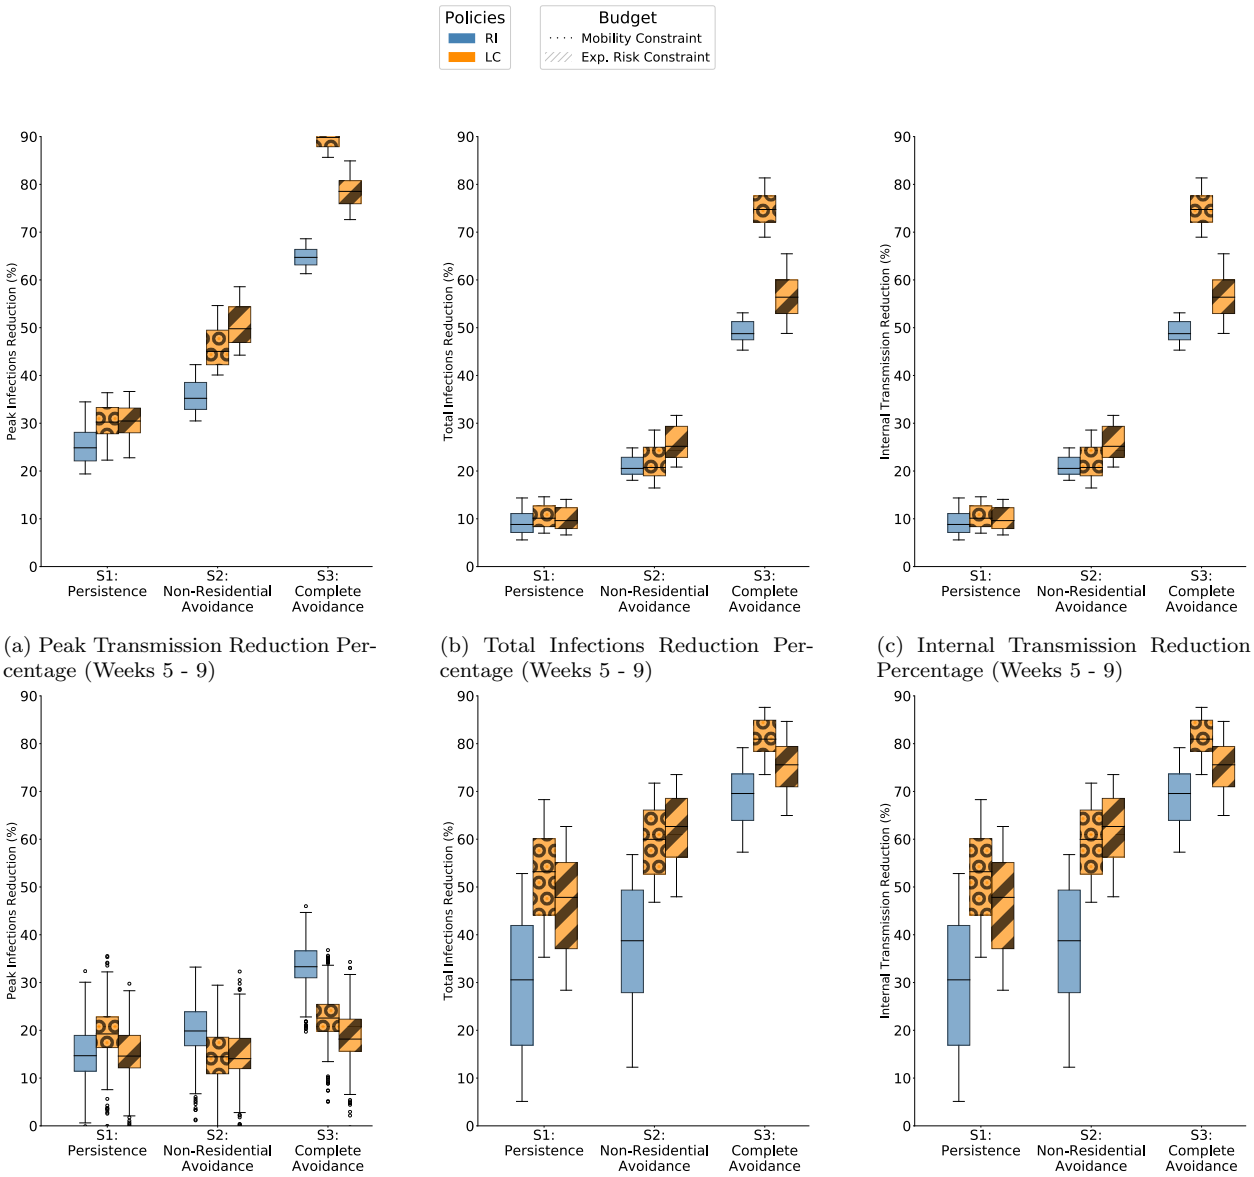

(a) Peak Transmission Reduction Percentage (Weeks 5 - 9)

(b) Total Infections Reduction Percentage (Weeks 5 - 9)

(c) Internal Transmission Reduction Percentage (Weeks 5 - 9)

(d) Peak Transmission Reduction Percentage (Weeks 10 - 14)

(e) Total Infections Reduction Percentage (Weeks 10 - 14)

(f) Internal Transmission Reduction Percentage (Weeks 10 - 14)

Figure S12: Disease control outcomes in Fall 2019 for $\mathsf{LC}_{\mathrm{PRank}}$. $(a-c)$ The ABM was calibrated on weeks $5-9$ of Fall 2020 at GT. Under all scenarios, for all outcomes, $\mathsf{LC}_{\mathrm{PRank}}$ is better than RI. $(d-f)$ The ABM was calibrated on weeks $10-14$ of Fall 2020 at GT. Under all scenarios, for all outcomes, $\mathsf{LC}_{\mathrm{PRank}}$ is better than RI.

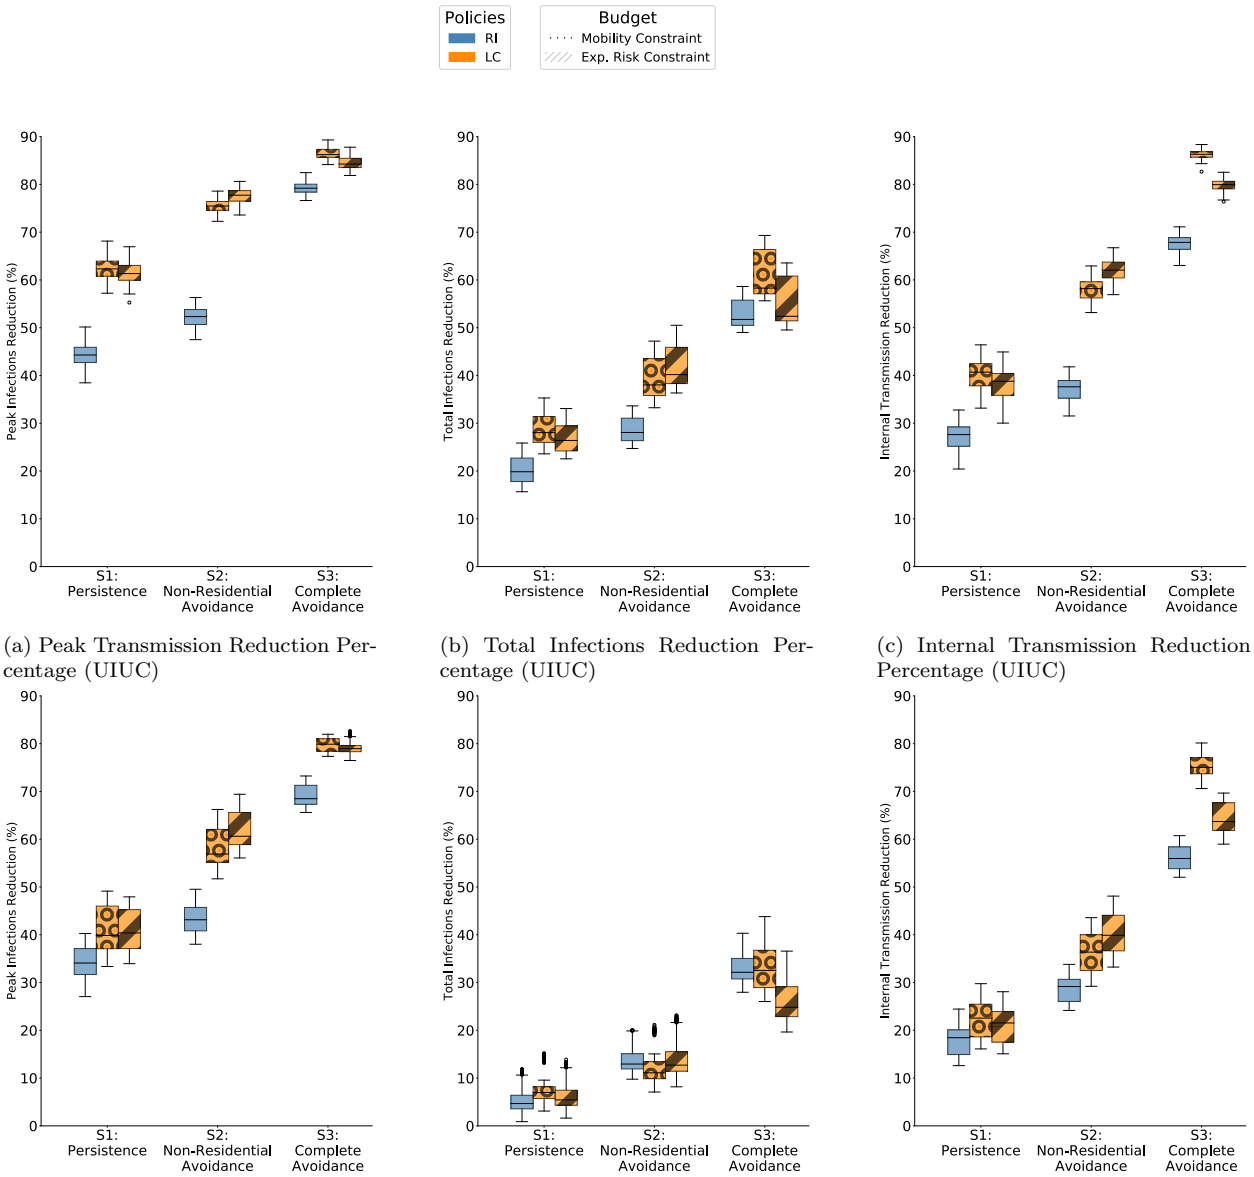

(a) Peak Transmission Reduction Percentage (UIUC)

(b) Total Infections Reduction Percentage (UIUC)

(c) Internal Transmission Reduction Percentage (UIUC)

(d) Peak Transmission Reduction Percentage (UCB)

(e) Total Infections Reduction Percentage (UCB)

(f) Internal Transmission Reduction Percentage (UCB)

Figure S13: Disease control outcomes in Fall 2019 for $\mathsf{LC}_{\text{PRank}}$. $(a-c)$ The ABM was calibrated on weeks $0-4$ of Fall 2020 at UIUC. Under all scenarios, for all outcomes, $\mathsf{LC}_{\text{PRank}}$ is better than RI. $(d-f)$ The ABM was calibrated on weeks $0-4$ of Fall 2020 at UC Berkeley. Under all scenarios, for all outcomes, $\mathsf{LC}_{\text{PRank}}$ is better than RI.

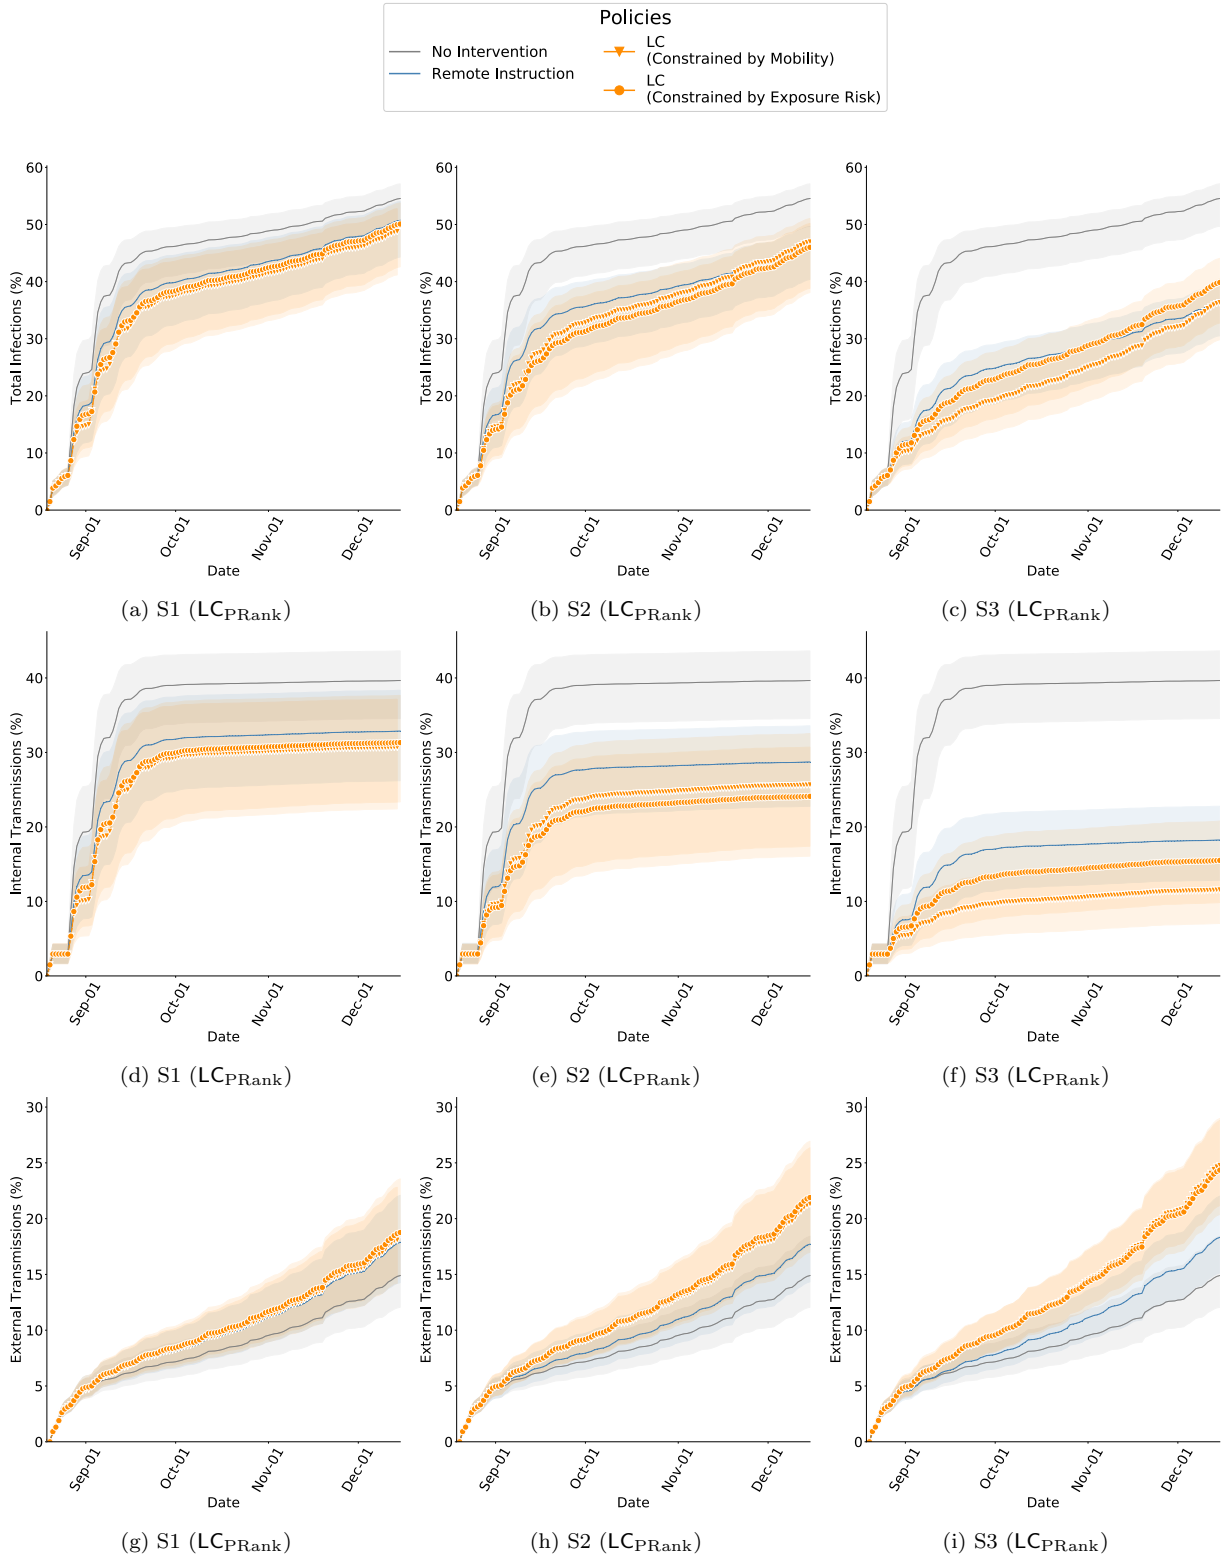

(a) S1 ($\mathsf{LC}_{\mathrm{PRank}}$) (b) S2 ($\mathsf{LC}_{\mathrm{PRank}}$) (c) S3 ($\mathsf{LC}_{\mathrm{PRank}}$)

(d) S1 ($\mathsf{LC}_{\mathrm{PRank}}$) (e) S2 ($\mathsf{LC}_{\mathrm{PRank}}$) (f) S3 ($\mathsf{LC}_{\mathrm{PRank}}$)

(g) S1 ($\mathsf{LC}_{\mathrm{PRank}}$) (h) S2 ($\mathsf{LC}_{\mathrm{PRank}}$) (i) S3 ($\mathsf{LC}_{\mathrm{PRank}}$)

Figure S14: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\mathrm{PRank}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\mathrm{PRank}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\mathrm{PRank}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

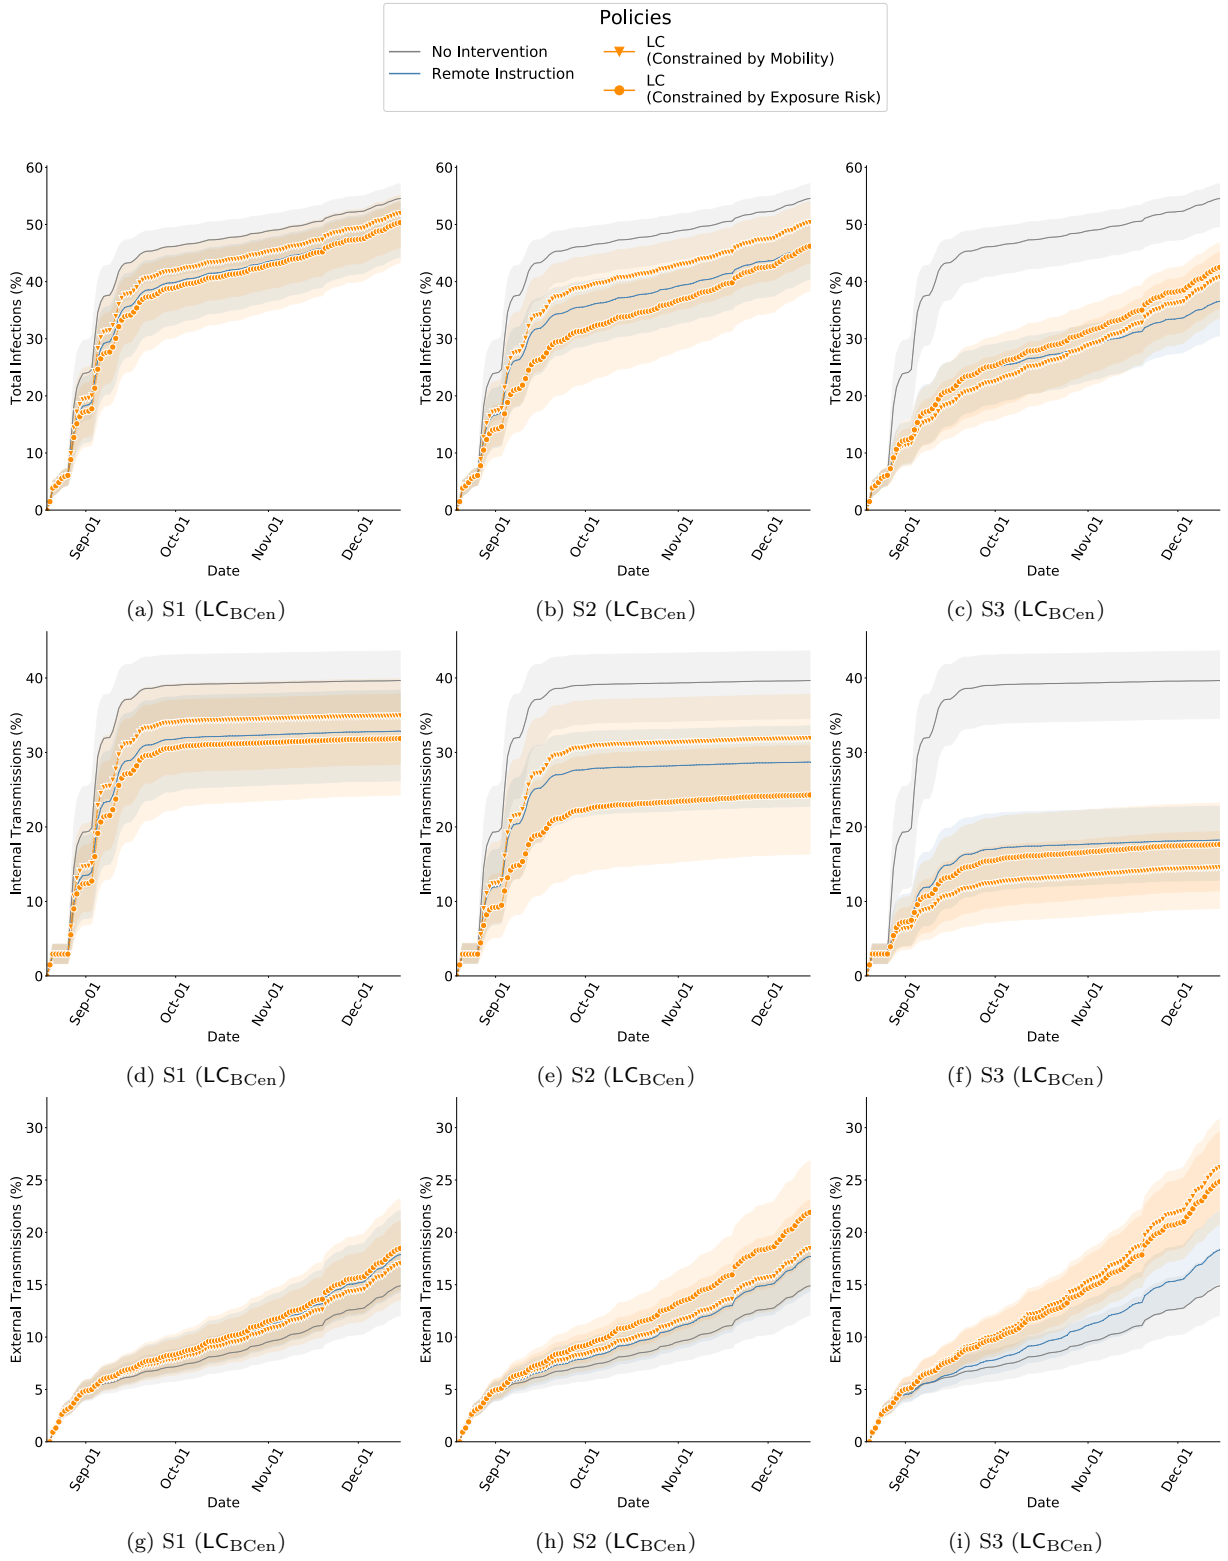

(a) S1 ($\mathsf{LC}_{\mathrm{BCen}}$) (b) S2 ($\mathsf{LC}_{\mathrm{BCen}}$) (c) S3 ($\mathsf{LC}_{\mathrm{BCen}}$)

(d) S1 ($\mathsf{LC}_{\mathrm{BCen}}$) (e) S2 ($\mathsf{LC}_{\mathrm{BCen}}$) (f) S3 ($\mathsf{LC}_{\mathrm{BCen}}$)

(g) S1 ($\mathsf{LC}_{\mathrm{BCen}}$) (h) S2 ($\mathsf{LC}_{\mathrm{BCen}}$) (i) S3 ($\mathsf{LC}_{\mathrm{BCen}}$)

Figure S15: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\mathrm{BCen}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\mathrm{BCen}}$ in comparison to RI, only when constrained under the exposure risk budget. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\mathrm{BCen}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

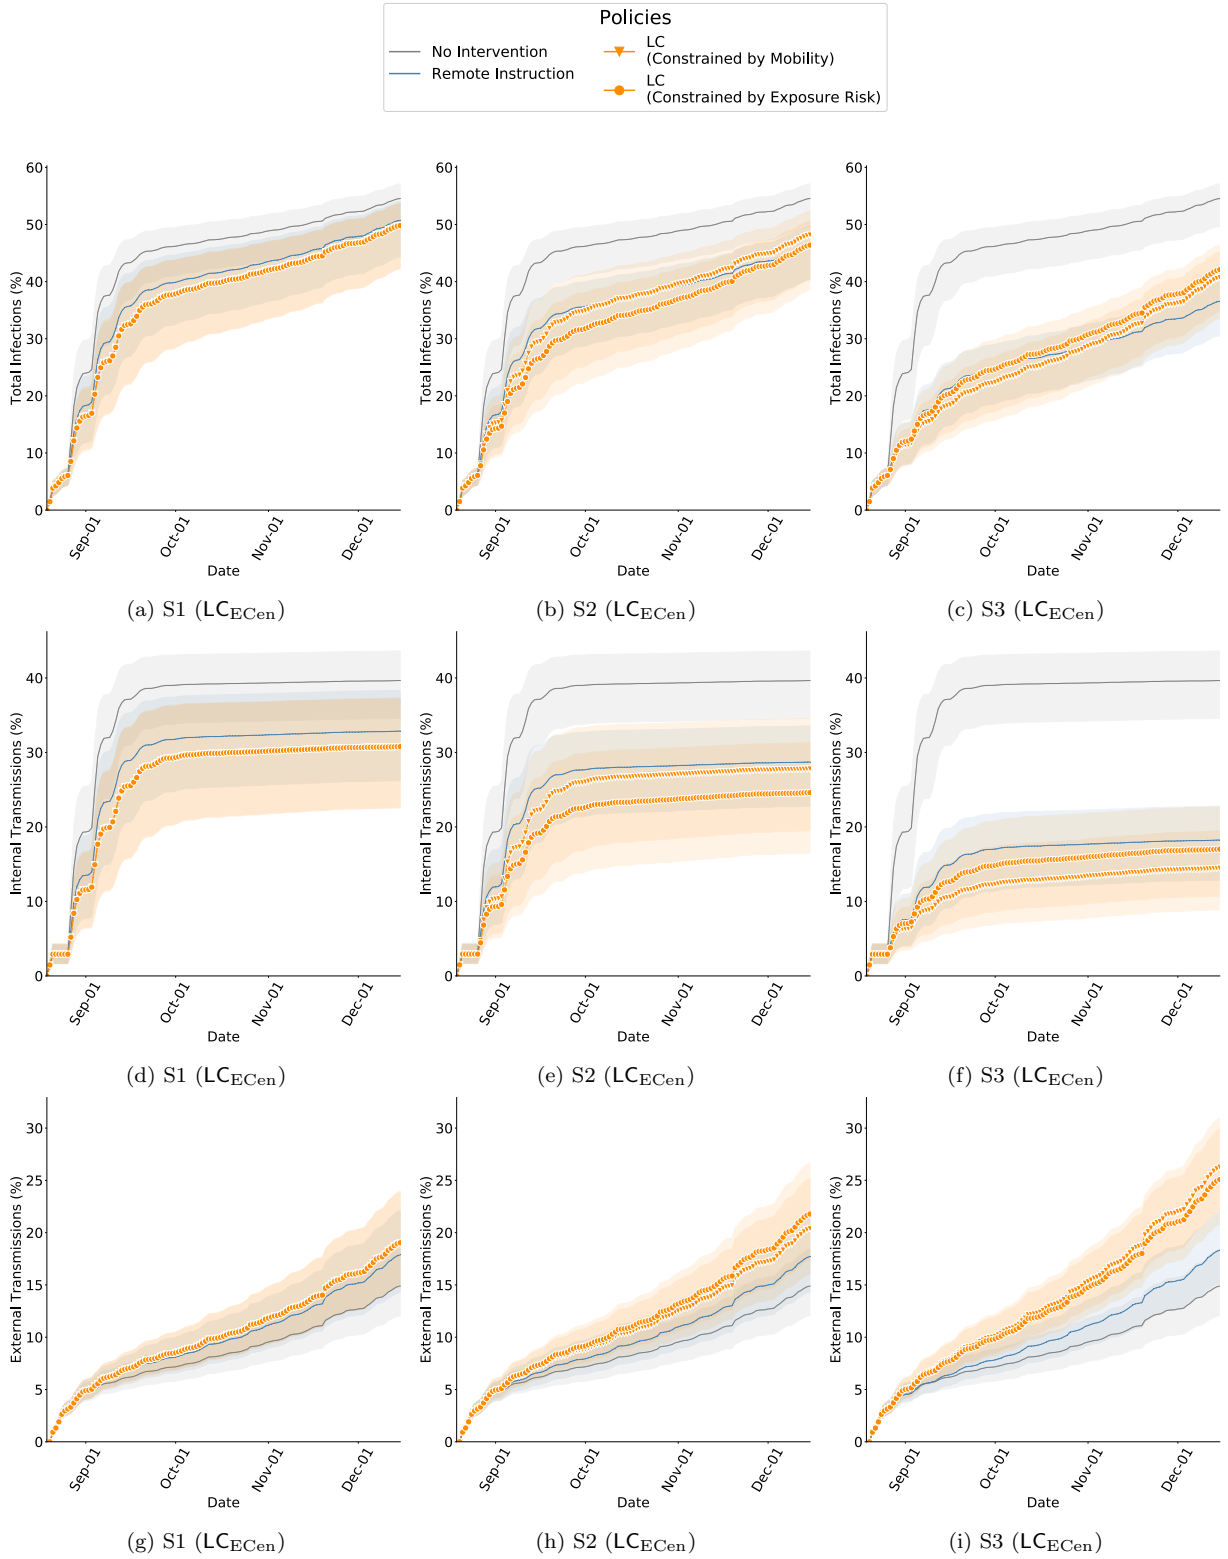

(a) S1 ($\mathsf{LC}_{\mathrm{ECen}}$) (b) S2 ($\mathsf{LC}_{\mathrm{ECen}}$) (c) S3 ($\mathsf{LC}_{\mathrm{ECen}}$)

(d) S1 ($\mathsf{LC}_{\mathrm{ECen}}$) (e) S2 ($\mathsf{LC}_{\mathrm{ECen}}$) (f) S3 ($\mathsf{LC}_{\mathrm{ECen}}$)

(g) S1 ($\mathsf{LC}_{\mathrm{ECen}}$) (h) S2 ($\mathsf{LC}_{\mathrm{ECen}}$) (i) S3 ($\mathsf{LC}_{\mathrm{ECen}}$)

Figure S16: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\mathrm{ECen}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\mathrm{ECen}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\mathrm{ECen}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

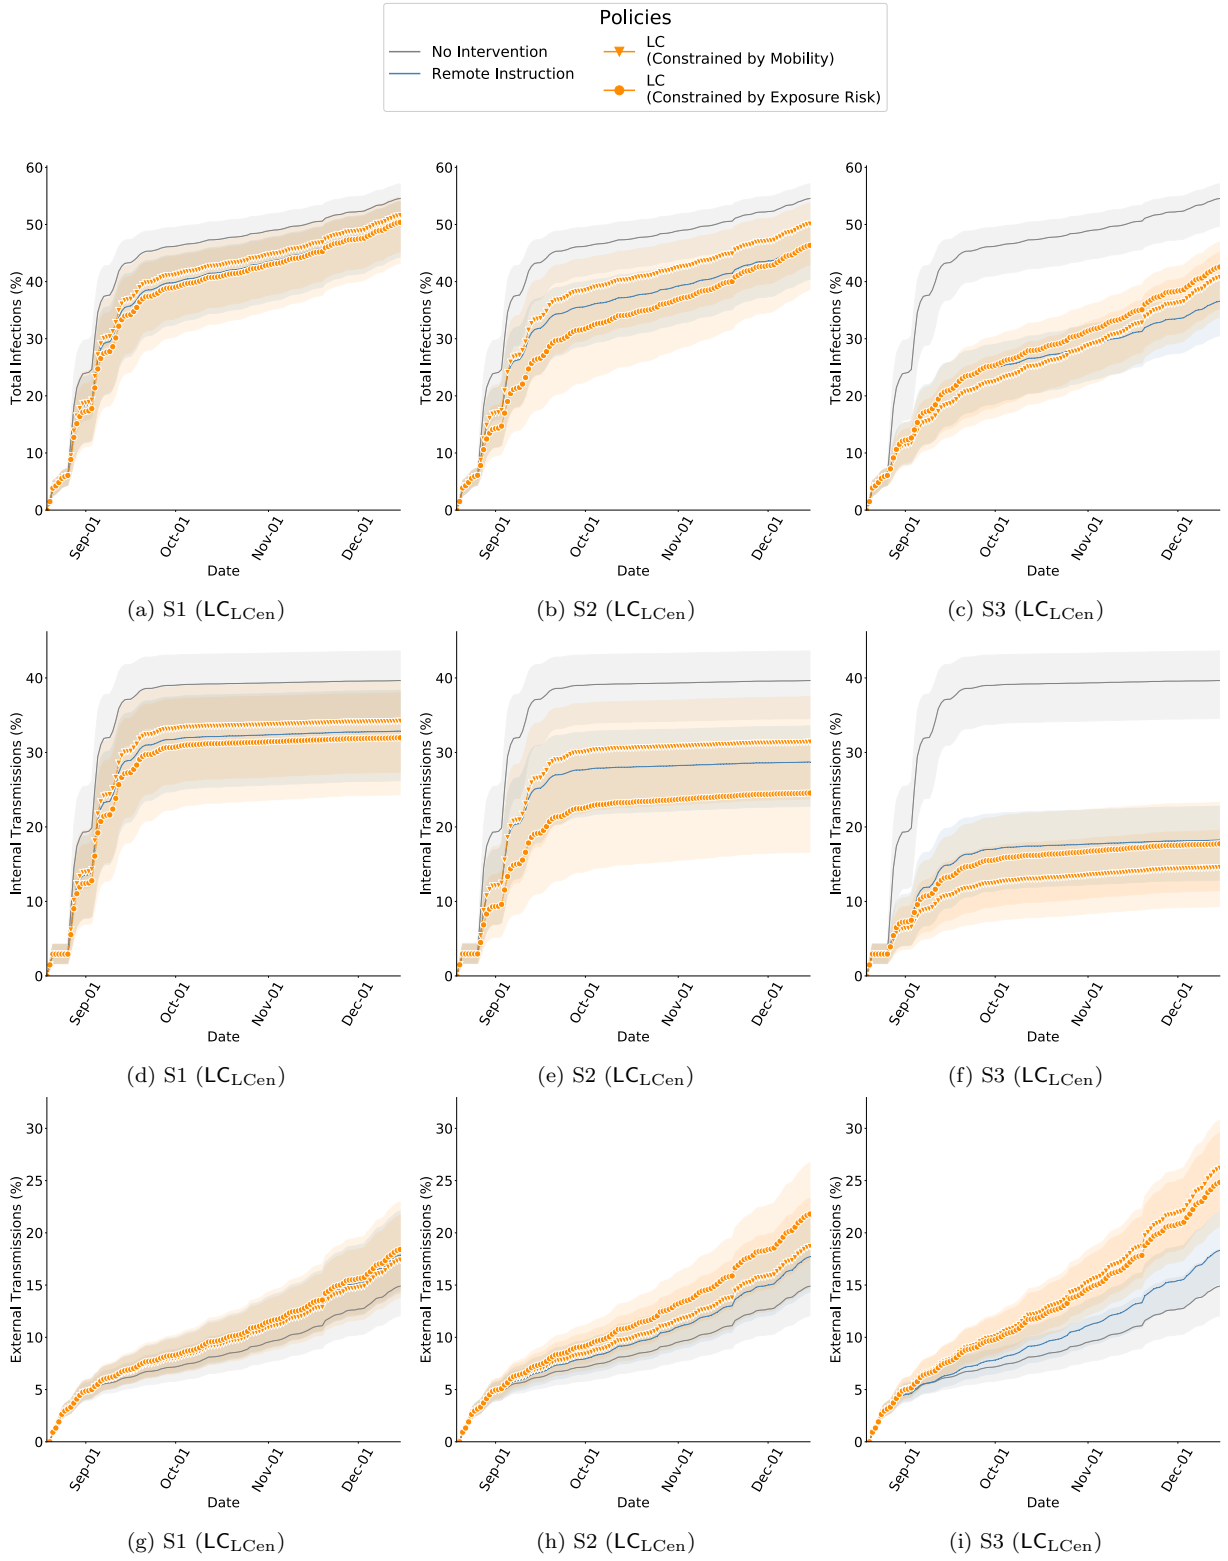

(a) S1 ($\mathsf{LC}_{\mathrm{LCen}}$) (b) S2 ($\mathsf{LC}_{\mathrm{LCen}}$) (c) S3 ($\mathsf{LC}_{\mathrm{LCen}}$)

(d) S1 ($\mathsf{LC}_{\mathrm{LCen}}$) (e) S2 ($\mathsf{LC}_{\mathrm{LCen}}$) (f) S3 ($\mathsf{LC}_{\mathrm{LCen}}$)

(g) S1 ($\mathsf{LC}_{\mathrm{LCen}}$) (h) S2 ($\mathsf{LC}_{\mathrm{LCen}}$) (i) S3 ($\mathsf{LC}_{\mathrm{LCen}}$)

Figure S17: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\mathrm{LCen}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\mathrm{LCen}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\mathrm{LCen}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

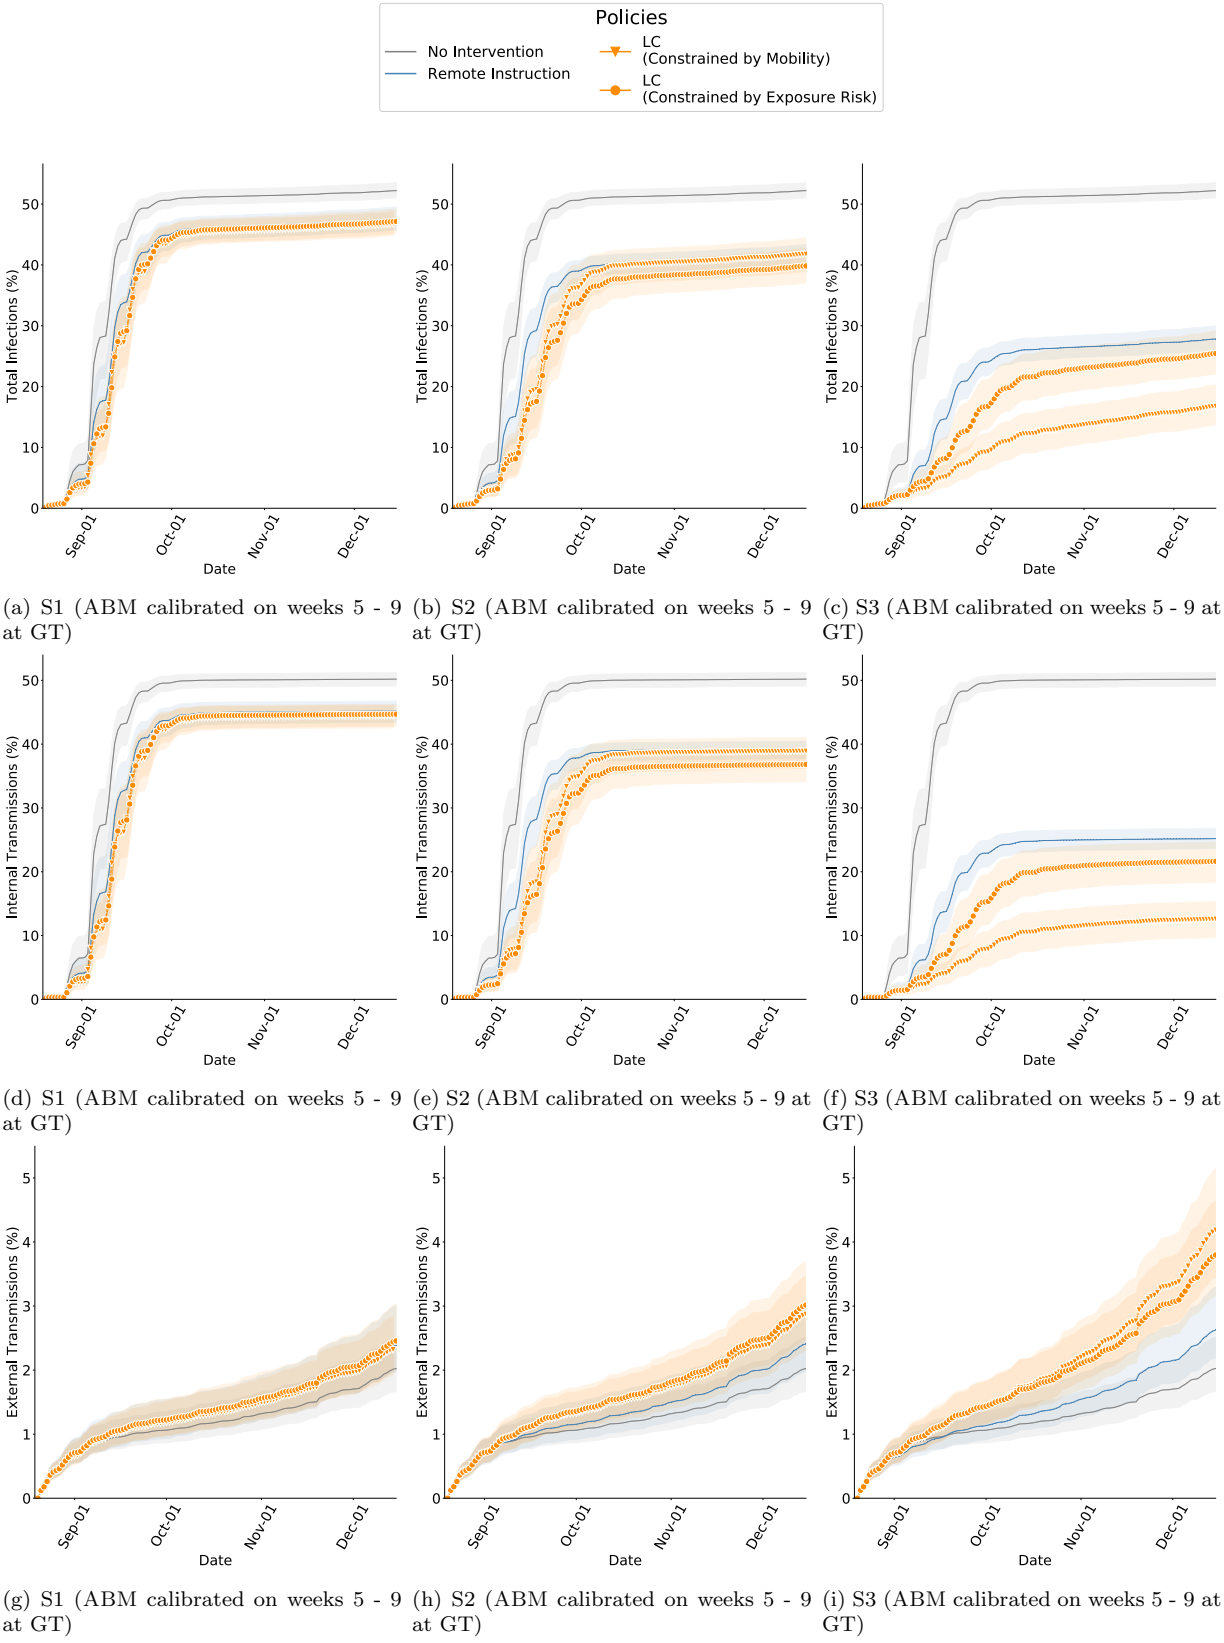

(a) S1 (ABM calibrated on weeks 5 - 9 at GT)

(b) S2 (ABM calibrated on weeks 5 - 9 at GT)

(c) S3 (ABM calibrated on weeks 5 - 9 at GT)

(d) S1 (ABM calibrated on weeks 5 - 9 at GT)

(e) S2 (ABM calibrated on weeks 5 - 9 at GT)

(f) S3 (ABM calibrated on weeks 5 - 9 at GT)

(g) S1 (ABM calibrated on weeks 5 - 9 at GT)

(h) S2 (ABM calibrated on weeks 5 - 9 at GT)

(i) S3 (ABM calibrated on weeks 5 - 9 at GT)

Figure S18: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\text{PRank}}$ with ABM calibrated on weeks $5-9$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

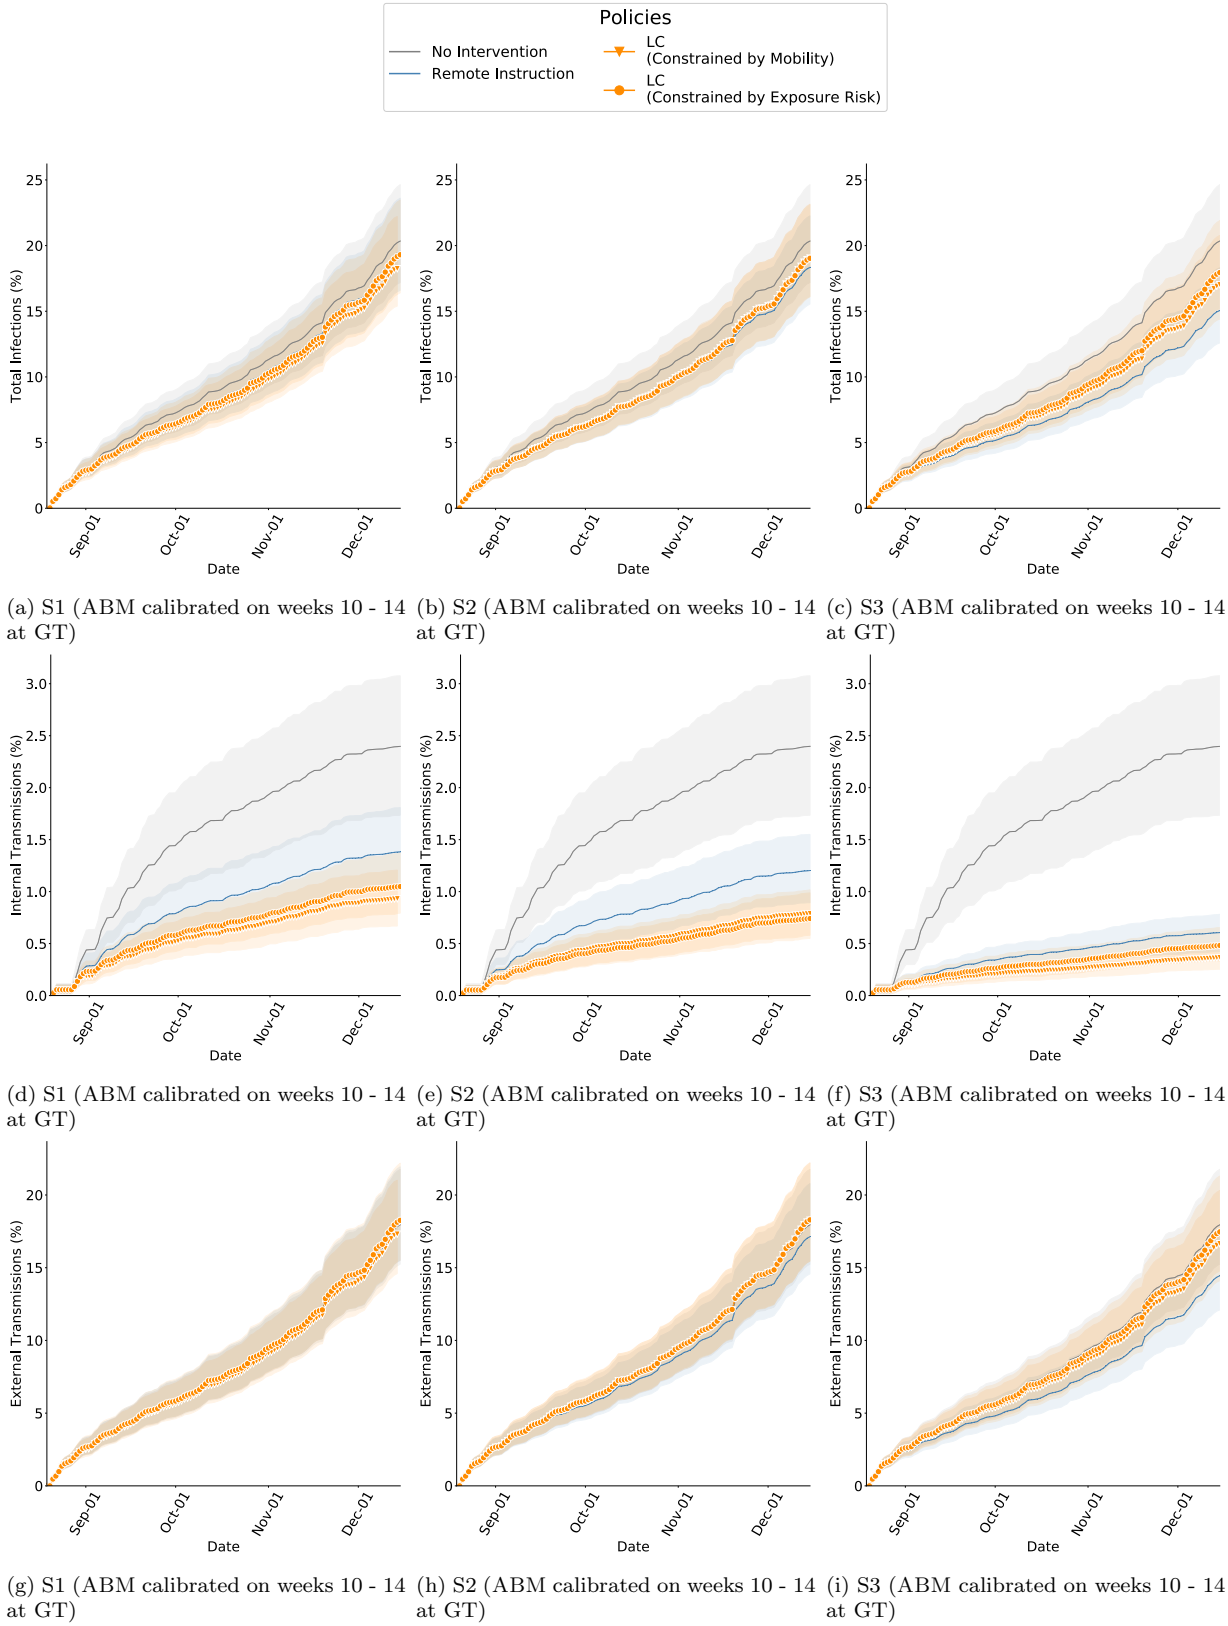

(a) S1 (ABM calibrated on weeks 10 - 14 at GT)

(b) S2 (ABM calibrated on weeks 10 - 14 at GT)

(c) S3 (ABM calibrated on weeks 10 - 14 at GT)

(d) S1 (ABM calibrated on weeks 10 - 14 at GT)

(e) S2 (ABM calibrated on weeks 10 - 14 at GT)

(f) S3 (ABM calibrated on weeks 10 - 14 at GT)

(g) S1 (ABM calibrated on weeks 10 - 14 at GT)

(h) S2 (ABM calibrated on weeks 10 - 14 at GT)

(i) S3 (ABM calibrated on weeks 10 - 14 at GT)

Figure S19: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\text{PRank}}$ with ABM calibrated on weeks $10-14$ of Fall 2020, GT. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

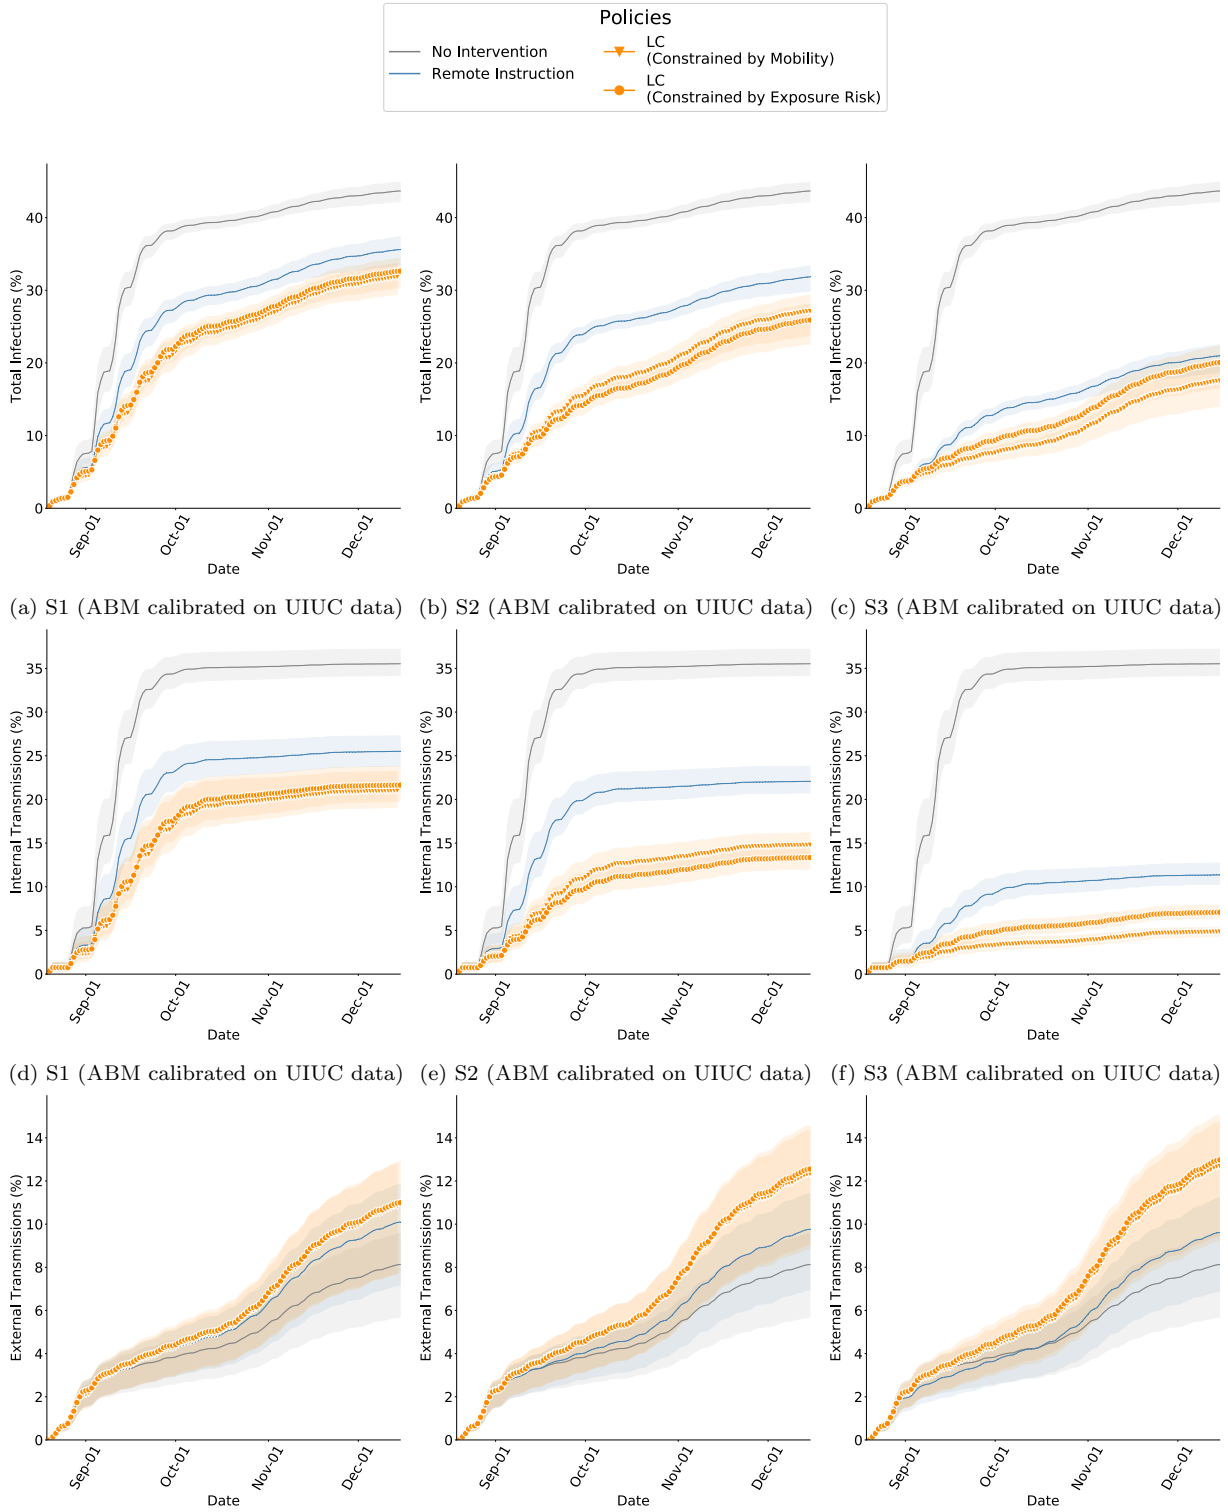

(a) S1 (ABM calibrated on UIUC data) (b) S2 (ABM calibrated on UIUC data) (c) S3 (ABM calibrated on UIUC data)

(d) S1 (ABM calibrated on UIUC data) (e) S2 (ABM calibrated on UIUC data) (f) S3 (ABM calibrated on UIUC data)

(g) S1 (ABM calibrated on UIUC data) (h) S2 (ABM calibrated on UIUC data) (i) S3 (ABM calibrated on UIUC data)

Figure S20: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\mathrm{PRank}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, UIUC. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\mathrm{PRank}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\mathrm{PRank}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

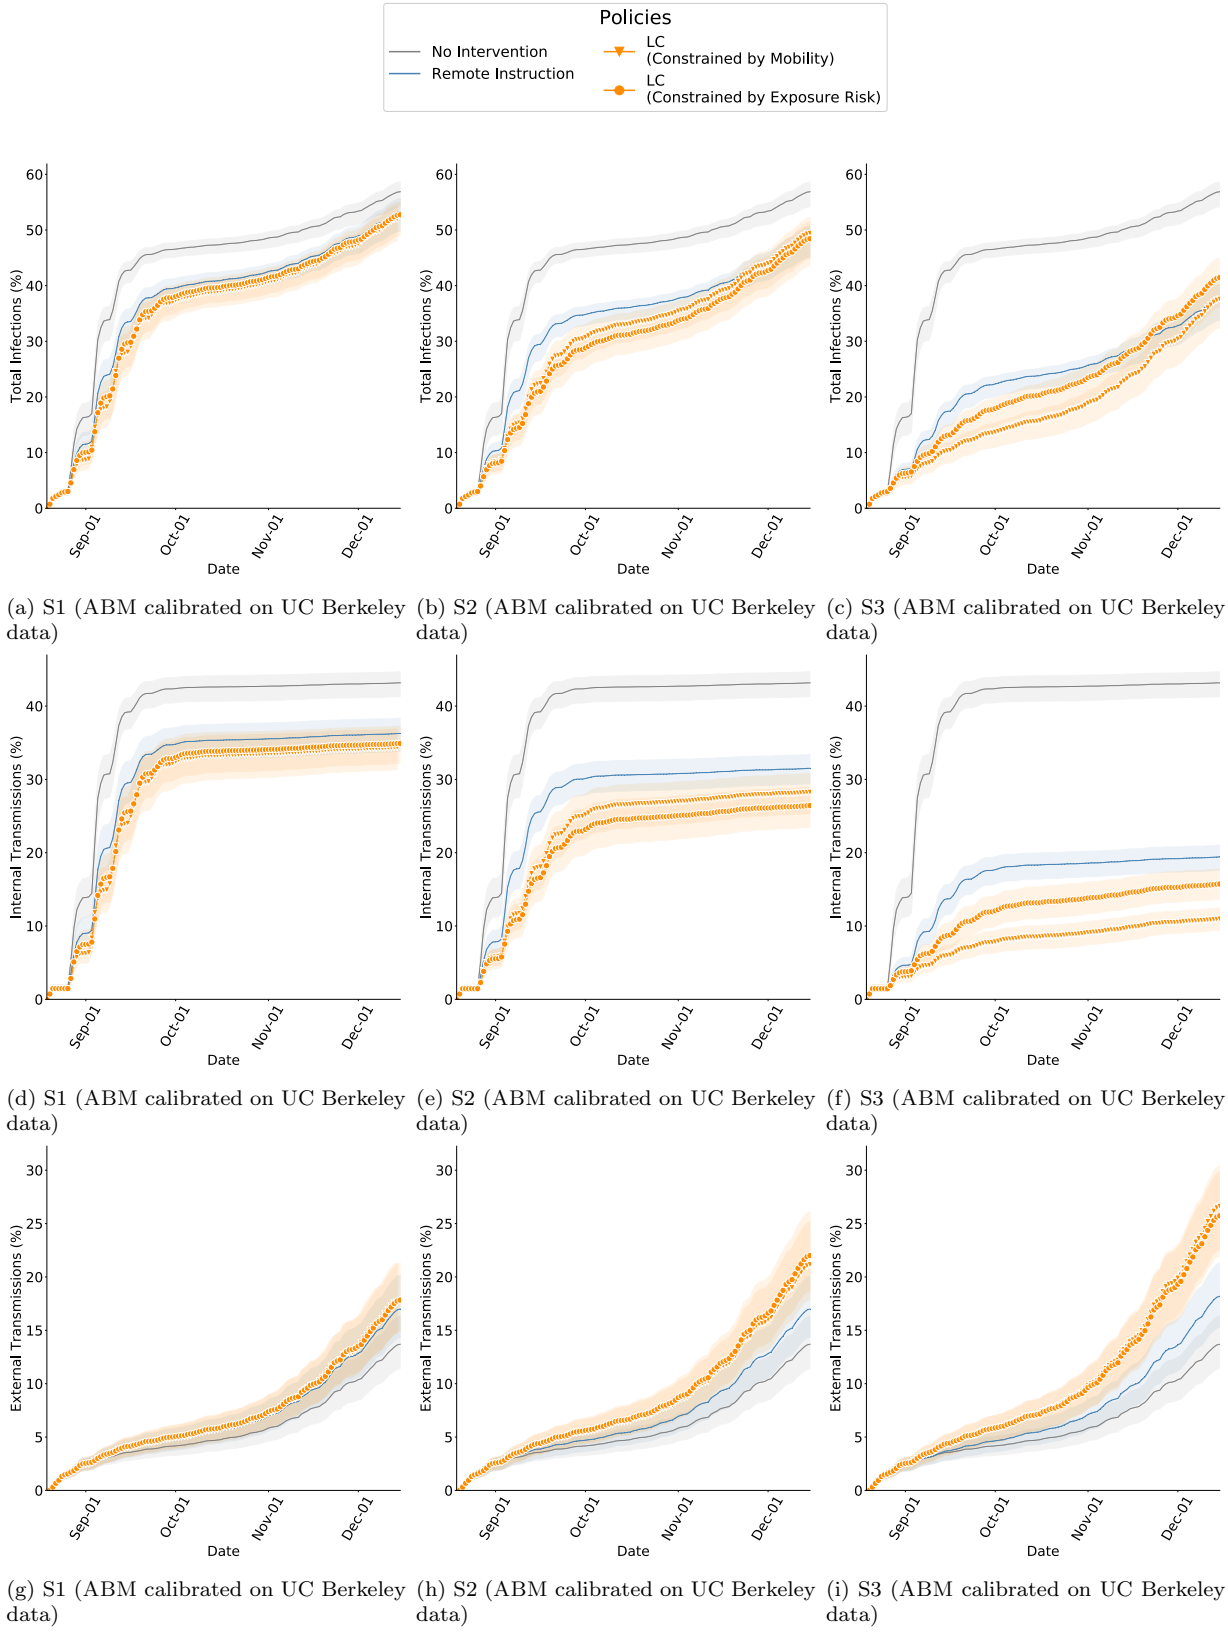

(a) S1 (ABM calibrated on UC Berkeley data)

(b) S2 (ABM calibrated on UC Berkeley data)

(c) S3 (ABM calibrated on UC Berkeley data)

(d) S1 (ABM calibrated on UC Berkeley data)

(e) S2 (ABM calibrated on UC Berkeley data)

(f) S3 (ABM calibrated on UC Berkeley data)

(g) S1 (ABM calibrated on UC Berkeley data)

(h) S2 (ABM calibrated on UC Berkeley data)

(i) S3 (ABM calibrated on UC Berkeley data)

Figure S21: Cumulative infections in Fall 2019 while comparing RI and $\mathsf{LC}_{\text{PRank}}$ with ABM calibrated on weeks $0-4$ of Fall 2020, UC Berkeley. The bands show the 2.75^{th} and 97.25^{th} percentile. $(a-c)$ Total infections of interventions is lower than no-intervention scenarios and is lowest in the *S3* scenario. In this scenario, the mobility budget is 69% of what it would be without interventions, and therefore the transmissions are also contained. In comparison, in Fall 2020, we saw far fewer infections which is because the mobility was 39% of that in Fall 2019. $(d-f)$ Internal transmissions are lower with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. $(g-i)$ External transmissions are higher with $\mathsf{LC}_{\text{PRank}}$ in comparison to RI. Since internal transmission is controlled, more individuals remain susceptible to infections from outside campus.

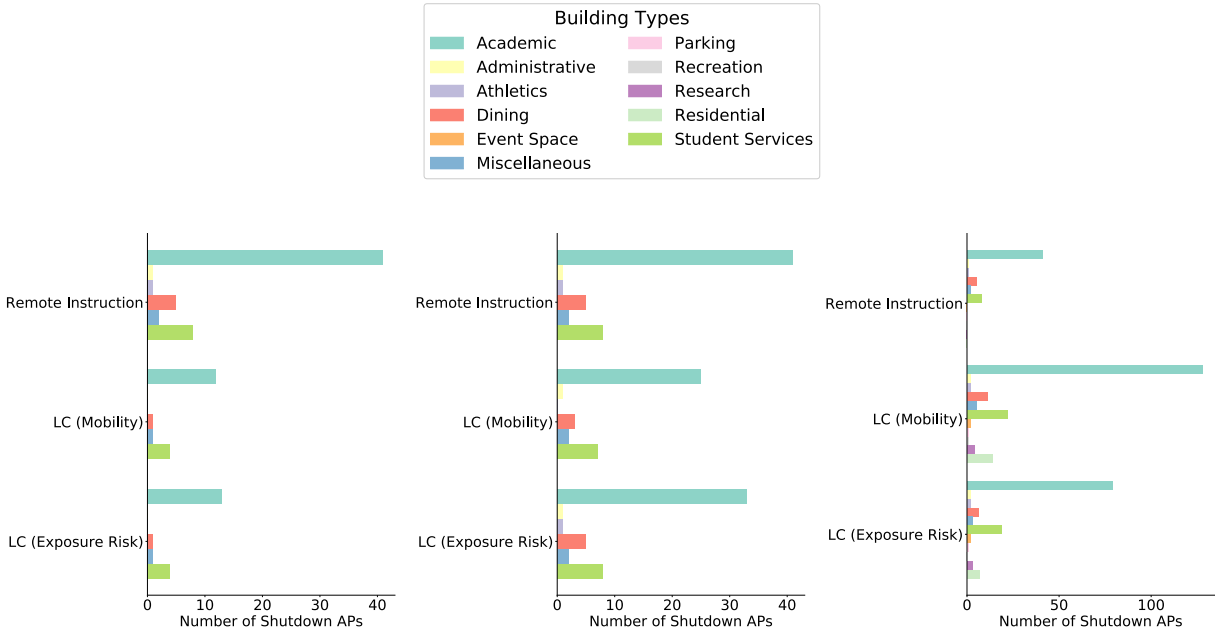

(a) S1 ($\mathsf{LC}_{\mathrm{PRank}}$) (b) S2 ($\mathsf{LC}_{\mathrm{PRank}}$) (c) S3 ($\mathsf{LC}_{\mathrm{PRank}}$)

Figure S22: The locations shutdown by each policy are grouped into the the general building category. The distribution of locations is different between policies, for example, in *S1* (*a*) and *S2* (*b*), LC closes fewer locations that RI. Even when targeting spaces in similar buildings, the locations are qualitatively different — RI only affects classrooms, whereas LC also closes smaller spaces like breakout rooms, reading areas and cafes. LC In *S3* (*c*) we find LC to target locations in a greater variety of buildings, but it also targets more locations to utilize the budget.

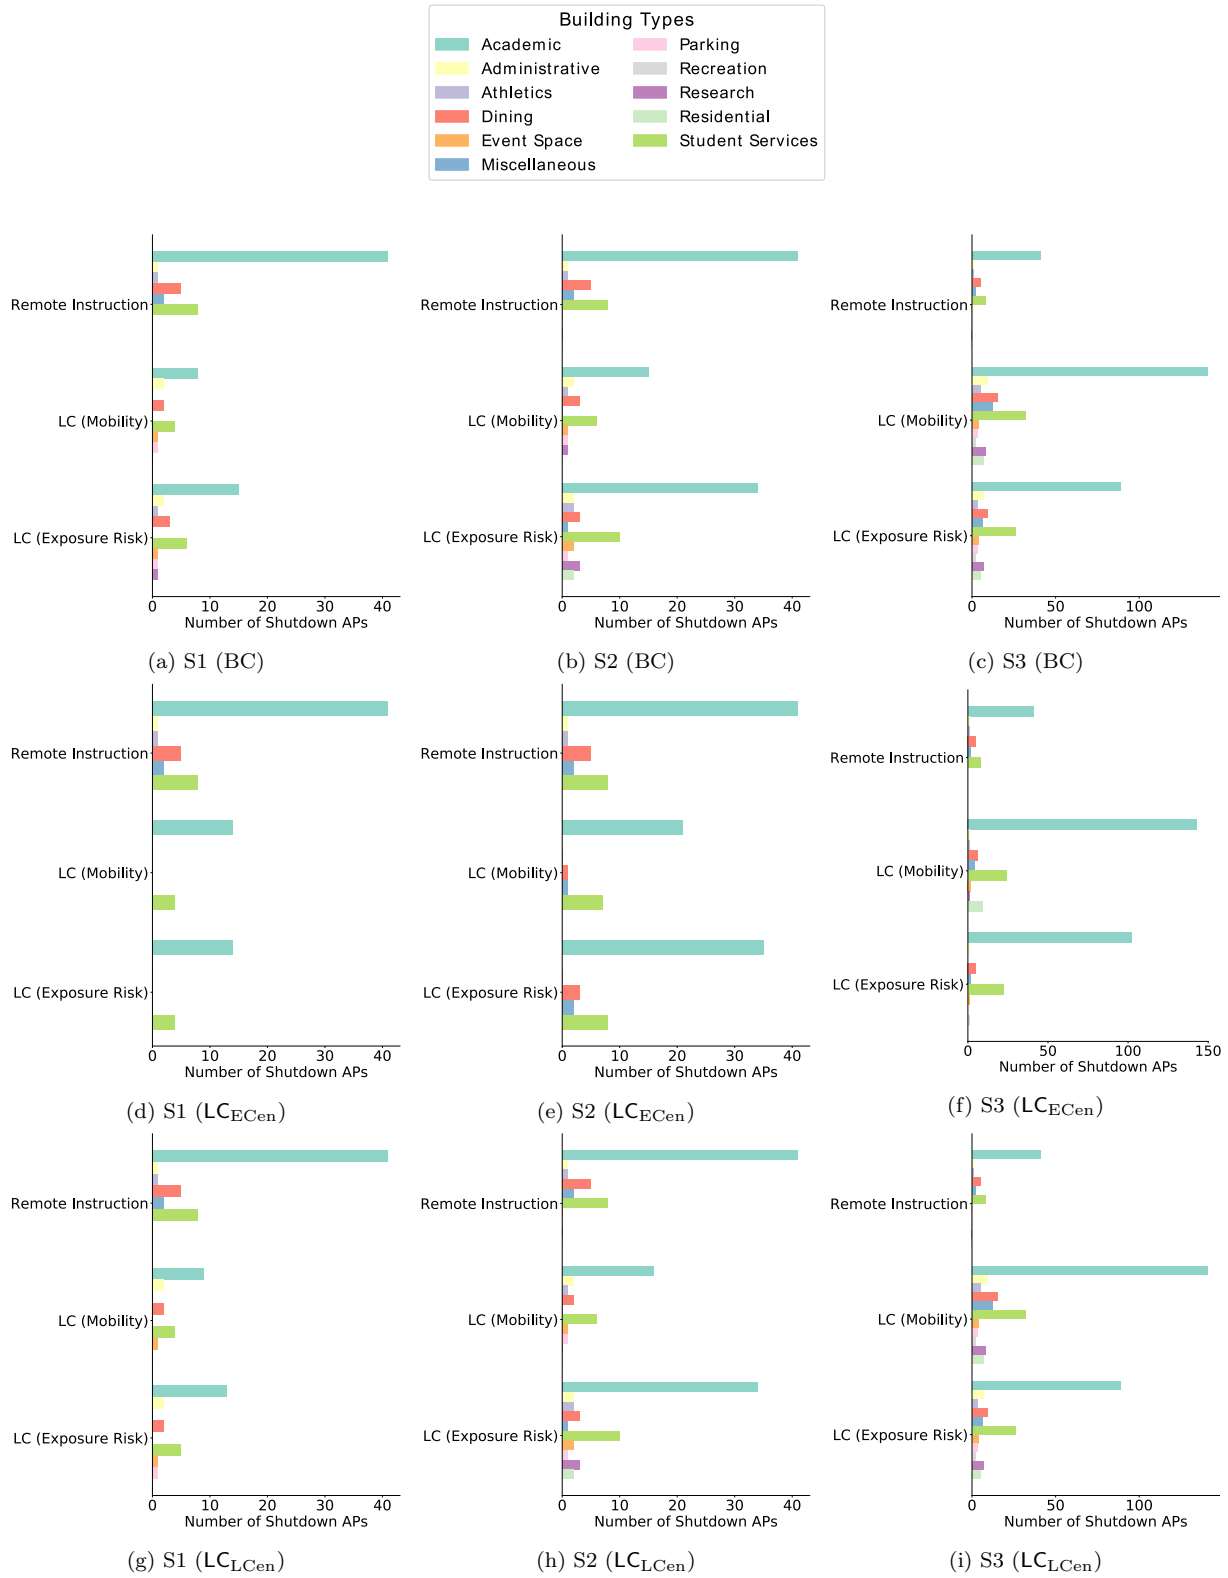

(a) S1 (BC) (b) S2 (BC) (c) S3 (BC)

(d) S1 ($\mathsf{LC}_{\mathrm{ECen}}$) (e) S2 ($\mathsf{LC}_{\mathrm{ECen}}$) (f) S3 ($\mathsf{LC}_{\mathrm{ECen}}$)

(g) S1 ($\mathsf{LC}_{\mathrm{LCen}}$) (h) S2 ($\mathsf{LC}_{\mathrm{LCen}}$) (i) S3 ($\mathsf{LC}_{\mathrm{LCen}}$)

Figure S23: The locations shutdown by each policy are grouped into the the general building category. The distribution of locations is different between policies, for example, in *S1* (*a*) and *S2* (*b*), LC closes fewer locations that RI. Even when targeting spaces in similar buildings, the locations are qualitatively different — RI only affects classrooms, whereas LC also closes smaller spaces like breakout rooms, reading areas and cafes. LC In *S3* (*c*) we find LC to target locations in a greater variety of buildings, but it also targets more locations to utilize the budget.

SI References

References from SI

[1] M. Andersen. Early evidence on social distancing in response to covid-19 in the united states. *Available at SSRN 3569368*, 2020.

[2] J. P. Azevedo, A. Hasan, D. Goldemberg, S. A. Iqbal, and K. Geven. *Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes: A set of global estimates*. The World Bank, 2020.

[3] E. Bagdasaryan, G. Berlstein, J. Waterman, E. Birrell, N. Foster, F. B. Schneider, and D. Estrin. Ancile: Enhancing privacy for ubiquitous computing with use-based privacy. In *Proceedings of the 18th ACM Workshop on Privacy in the Electronic Society*, pages 111–124, 2019.

[4] S. G. Benzell, A. Collis, and C. Nicolaides. Rationing social contact during the covid-19 pandemic: Transmission risk and social benefits of us locations. *Proceedings of the National Academy of Sciences*, 117(26):14642–14644, 2020.

[5] P. Bonacich. Some unique properties of eigenvector centrality. *Social networks*, 29(4):555–564, 2007.

[6] M. Borowiak, F. Ning, J. Pei, S. Zhao, H.-R. Tung, and R. Durrett. Controlling the spread of covid-19 on college campuses. *arXiv preprint arXiv:2008.07293*, 2020.

[7] G. A. Bruneau, V. C. Müller, and M. S. Gilthorpe. The ethical imperatives of the covid 19 pandemic: A review from data ethics. *Veritas: Revista de Filosofía y Teología*, 46:13–35, 2020.

[8] S. Chang, E. Pierson, P. W. Koh, J. Gerardin, B. Redbird, D. Grusky, and J. Leskovec. Mobility network models of covid-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, 2021.

[9] Cisco. Wi-fi location-based services 4.1 design guide white paper, 2014. Available at: https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Mobility/WiFiLBS-DG/wifich2.html.

[10] V. Das Swain, H. Kwon, B. Saket, M. B. Morshed, K. Tran, D. Patel, Y. Tian, J. Philipose, Y. Cui, T. Plötz, et al. Leveraging wifi network logs to infer social interactions: A case study of academic performance and student behavior. *arXiv e-prints*, 2020.

[11] A. C. P. H. Department. Real-time data of the impact of covid-19, 2020. Available at: https://covid-19.acgov.org/data.page.

[12] C.-U. P. H. Distric. Champaign-urbana covid-19 coronavirus information, 2020. Available at: https://www.c-uphd.org/champaign-urbana-illinois-coronavirus-information.html.

[13] E. Dorn, B. Hancock, J. Sarakatsannis, and E. Viruleg. Covid-19 and student learning in the united states: The hurt could last a lifetime. *McKinsey & Company*, 2020.

[14] M. H. S. Eldaw, M. Levene, and G. Roussos. Presence analytics: making sense of human social presence within a learning environment. In *2018 IEEE/ACM 5th International Conference on Big Data Computing Applications and Technologies (BDCAT)*, pages 174–183. IEEE, 2018.

[15] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.

[16] K. Gaythorpe, N. Imai, G. Cuomo-Dannenburg, M. Baguelin, S. Bhatia, A. Boonyasiri, and A. Cori. Report 8: Symptom progression of covid-19, 2020.

[17] G. Gibson, J. S. Weitz, M. P. Shannon, B. Holton, A. Bryksin, B. Liu, S. Bramblett, J. Williamson, M. Farrell, A. Ortiz, C. T. Abdallah, and A. J. García. Surveillance-to-diagnostic testing program for asymptomatic sars-cov-2 infections on a large, urban campus - georgia institute of technology, fall 2020. *medRxiv*, 2021. Available at: https://www.medrxiv.org/content/early/2021/01/31/2021.01.28.21250700.

[18] P. T. Gressman and J. R. Peck. Simulating covid-19 in a university environment. *Mathematical biosciences*, 328:108436, 2020.

[19] E. S. Gurley. Strategies to support the covid-19 response in lmics, 2020. Available at: https://hopkinsglobalhealth.org/assets/documents/CGH_Webinar_-_Contact_Tracing_(Final_Version).pdf.

[20] P. Keskinocak, B. E. Oruc, A. Baxter, J. Asplund, and N. Serban. The impact of social distancing on covid19 spread: State of georgia case study. *Plos one*, 15(10):e0239798, 2020.

[21] W. H. Kruskal and W. A. Wallis. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621, 1952.

[22] B. Lopman, C. Liu, A. Le Guillou, A. Handel, T. L. Lash, A. Isakov, and S. Jenness. A model of covid-19 transmission and control on university campuses. *medRxiv*, 2020.

[23] C. Makridis and J. Hartley. The cost of covid-19: A rough estimate of the 2020 us gdp impact, 2020.

[24] D. Mangrum and P. Niekamp. Jue insight: College student travel contributed to local covid-19 spread. *Journal of Urban Economics*, page 103311, 2020.

[25] K. I. McKinnon. Convergence of the nelder–mead simplex method to a nonstationary point. *SIAM Journal on optimization*, 9(1):148–158, 1998.

[26] S. M. Moghadas, M. C. Fitzpatrick, P. Sah, A. Pandey, A. Shoukat, B. H. Singer, and A. P. Galvani. The implications of silent transmission for the control of covid-19 outbreaks. *Proceedings of the National Academy of Sciences*, 117(30):17513–17515, 2020.

[27] A. W.-F. L. Monitor. Wifi location monitor accuracy, 2016. Available at: https://www.accuware.com/support/wi-fi-location-monitor-accuracy/.

[28] S. J. Mooney and V. Pejaver. Big data in public health: terminology, machine learning, and privacy. *Annual review of public health*, 39:95–112, 2018.

[29] M. E. Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical review E*, 64(1):016132, 2001.

[30] A. Nierenberg and A. Pasick. Schools briefing: University outbreaks and parental angst, 2020. Available at: https://www.nytimes.com/2020/08/19/us/colleges-closing-covid.html.

[31] U. of Illinois at Urbana-Champaign. On-campus covid-19 testing, 2020. Available at: https://covid19.illinois.edu/on-campus-covid-19-testing-data-dashboard/.

[32] G. D. of Public Health. Georgia department of public health daily status report, 2020. Available at: https://dph.georgia.gov/covid-19-daily-status-report.

[33] G. I. of Technology. Georgia tech launches campus coronavirus testing, 2020. Available at: https://health.gatech.edu/coronavirus/testing-launched.

[34] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.

[35] B. Pfefferbaum and C. S. North. Mental health and the covid-19 pandemic. *New England Journal of Medicine*, 383(6):510–512, 2020.

[36] A. Pfitzmann and M. Hansen. A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, 2010.

[37] J. H. Stone, M. J. Frigault, N. J. Serling-Boyd, A. D. Fernandes, L. Harvey, A. S. Foulkes, N. K. Horick, B. C. Healy, R. Shah, A. M. Bensaci, et al. Efficacy of tocilizumab in patients hospitalized with covid-19. *New England Journal of Medicine*, 383(24):2333–2344, 2020.

[38] B. University of California. Coronavirus dashboard testing, 2020. Available at: https://coronavirus.berkeley.edu/dashboard/.

[39] S. Venkatramanan, A. Sadilek, A. Fadikar, C. L. Barrett, M. Biggerstaff, J. Chen, X. Dotiwalla, P. Eastham, B. Gipson, D. Higdon, et al. Forecasting influenza activity using machine-learned mobility map. *Nature Communications*, 12(1):1–12, 2021.

[40] J. L. Wang and M. C. Loui. Privacy and ethical issues in location-based tracking systems. In *2009 IEEE International Symposium on Technology and Society*, pages 1–4. IEEE, 2009.

[41] S. Ware, C. Yue, R. Morillo, J. Lu, C. Shang, J. Kamath, A. Bamis, J. Bi, A. Russell, and B. Wang. Large-scale automatic depression screening using meta-data from wifi infrastructure. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(4):1–27, 2018.

[42] S. Watson, S. Hubler, D. Ivory, and R. Gebeloff. A new front in america's pandemic: College towns, 2020. Available at: https://www.nytimes.com/2020/09/06/us/colleges-coronavirus-students.html.

[43] K. A. Weeden and B. Cornwell. The small-world network of college classes: implications for epidemic spread on a university campus. *Sociological science*, 7:222–241, 2020.

[44] B. Wilder, M. Charpignon, J. A. Killian, H.-C. Ou, A. Mate, S. Jabbari, A. Perrault, A. N. Desai, M. Tambe, and M. S. Majumder. Modeling between-population variation in covid-19 dynamics in hubei, lombardy, and new york city. *Proceedings of the National Academy of Sciences*, 117(41):25904–25910, 2020.

[45] D. Wu, H. Hamilton, L. Jagrowski, D. Nazzal, and L. N. Steimle. Revisiting the small world of college classes: The trade-offs between risk and burden of partial remote instruction, 2021. Working paper.