Opening of hotels and ski facilities: impact on mobility, spending, and Covid-19 outcomes

Abstract

This paper investigates how the opening of hotels and ski facilities in Poland impacted touristic spending, mobility and Covid-19 outcomes. We use administrative data from a government program subsidizing travel to show that the policy increased consumption of touristic services in ski resorts. Next, leveraging geolocation data from Facebook, we show that ski resorts experienced a significant influx of tourists, increasing the number of local users by up to 50%. Furthermore, we show that there was an increase in the probability of meetings between pairs of users from distanced locations and pairs of users from touristic and non-touristic areas. As the policy impacted travels and gatherings, we then analyze its effect on the diffusion of Covid-19. We find a significant association between touristic movements and the severity of a major pandemic wave in Poland. In particular, we observe that counties with ski facilities experienced more infections after the opening. Moreover, counties strongly connected to the ski resorts during the opening had more subsequent cases than weakly connected counties.

Keywords: Covid-19, Tourism, Risky behaviors, Mobility, Spatial Spillovers

1 Introduction

Regulating risky behaviors often requires balancing competing policy goals. Maximizing utility or income through risky activity comes at the expense of the health of risk-takers and individuals not engaged in the risky behavior. Such negative externalities present a particular challenge in designing an efficient policy, as decision-makers often lack information on the magnitude of the social costs of risky behaviors.

The covid-19 pandemic offers a unique setting to investigate the trade-off between individual freedoms and the negative externalities they generate (Stiglitz [2021], Stoddard et al. [2021]). This study opportunity arises as authorities attempted to balance stimulating economic activity and preventing infections (Alvarez et al. [2020], Acemoglu et al. [2021], Caulkins et al. [2021]).

The opening of tourism is a particularly useful case. Engaging in tourism has a large potential for negative spillovers during a pandemic. Tourism creates long-distance movements of the population (Mangrum and Niekamp [2020]), and, as such, can contribute to the spread of infectious diseases (Belik et al. [2011], Bajardi et al. [2011], Findlater and Bogoch [2018]). Simultaneously, tourists generate significant income for local economies. Hence, any decisions concerning tourist activity require balancing the trade-off between economic and public health outcomes. Yet, there is currently no evidence quantifying the impact of touristic mobility on the diffusion of infections. This paper aims to fill this gap by analyzing the effects of the reopening of Polish tourism on tourist consumption, mobility, and the spread of Covid-19.

Our study design takes advantage of a unique policy that reopened all hotels and ski lifts in Poland. On the 12th of February 2021, the Polish government reopened ski lifts and hotels at 50% capacity and with food supplied through room service only. At the same time, authorities reopened only cinemas, theatres, and operas at 50% capacity with mandatory masks[1]. The hotels and ski lifts remained open until the 20th of March, when the second wave of infections ravaged the country.

This setting is particularly relevant as ski resorts play an important role in the local economies and in the Polish tourism industry. While on the national scale only 2.2% of working Poles are employed in businesses

[1]Such entertainment facilities rarely operate in ski resorts

related to hospitality or recreation (H&R), areas with ski resorts rely heavily on income from these sectors. Almost 10% of all businesses active in towns with more than 3 ski lifts are related to H&R [2]. These locations are vital for the tourism industry. 9% of country's H&R businesses were located in communes with some ski facilities, which constitute only 3% of all communes. Moreover, these communes are home to 15.5% of all hotel beds in the country. Hence, the policy of opening hotels and ski resorts had a potential to strongly affect economic activities related to tourism.

We show that the policy indeed caused a significant increase in touristic spending and movements of tourists. First, we rely on data on the usage of government travel subsidies to document a swift growth in the consumption of touristic services in ski resorts after the opening. Second, using aggregated and anonymized geolocation data from Facebook (FB) and an event study framework (similar to Dave et al. [2021a]), we show that the policy's implementation increased the number of FB users in ski resorts by 25%-50%. Moreover, there has been a surge in the probability that users from non-touristic and touristic areas meet and in the probability that users living far away meet.

While the number of travels and tourist gatherings increased due to the policy, its effect on Covid-19 outcomes is a priori not obvious. On the one hand, visitors could carry the disease from their homes to touristic locations and back. On the other hand, they could only gather in their rooms or outdoors as restaurants were closed. To learn more about the impact of the policy on Covid-19 cases, we leverage granular infection data and the uncofoundedness approach from Callaway and Li [2021]. We use this newly developed method because traditional fixed effect models perform poorly in the presence of non-linearities inherent to patterns of Covid-19 diffusion (Gauthier [2021], Goodman-Bacon and Marcus [2020]). We find that counties with ski resorts see additional cases already in the third week after the opening. Such an early effect is absent in other counties. Moreover, counties strongly interacting with ski resorts during the policy have a higher incidence of infections than counties with weak interactions. This is consistent with a secondary spread from tourists bringing the virus back home.

The epidemic has sparked a large body of literature related to Covid-19, however relatively little research has been dedicated to the effect of tourism on infections. Researchers have shown that industry closures and stay-at-home orders have a limiting impact on both mobility and subsequent Covid-19 outcomes (Fang et al. [2020], Gupta et al. [2020], Lyu and Wehby [2020], Beria and Lunkar [2021], Courtemanche et al. [2020], Abouk and Heydari [2020], Badr et al. [2020], Lau et al. [2020], Morley et al. [2020], Xu et al. [2020], Goolsbee and Syverson [2021]). Fewer papers analyzed reopening. An exception is Nguyen et al. [2020] who found that lifting restrictions led to 6-8% increase in mobility.

Opening the tourism industry can lead to travels and large gatherings, and there has been some evidence relating these phenomena to the viral spread. One of the seminal papers on this topic is Adda [2016], which shows that school vacation and transportation strikes disrupt viral transmission. More recently, Chernozhukov et al. [2021], Andersen et al. [2021], Courtemanche et al. [2021], Bravata et al. [2021] , Goldhaber et al. [2021] link school and colleges operating modality to the local prevalence of infections. However, the opening of schools can have substantially different effects than tourism opening as children are less likely to suffer severe consequences of Covid-19 (Castagnoli et al. [2020], Dong et al. [2020]).

Closer to our population of interest (adults) are studies analyzing sport, social, and political gatherings. Large sporting events such as hockey, basketball and football game can lead to higher Covid-19 prevalence (Carlin et al. [2021], Alexander et al. [2020], Breidenbach and Mitze [2021]). Similarly, smaller gathering such as birthdays or bar meetings increase likelihood of subsequent infections (Harris [2020], Whaley et al. [2021]). The evidence in the case of political gatherings is mixed. Palguta et al. [2021] find an increase in the growth rate of Covid-19 in areas with elections, while Dave et al. [2020] conclude that a political rally in Tulsa did not affect local Covid cases. They note, however, that the local population enhanced social distancing, which could offset the effect of the gathering. This compensatory behavior is unlikely to have taken place for tourism opening because local population populations will interact with incoming tourists by providing them hospitality services. Economic and educational gatherings are also conducive to enhance local diffusion of infections. Taylor et al. [2020] show that proximity to livestock plant is associated with

[2]According to the data from December 2019 from the National Register of Business Entities (REGON). On the national level only 3% of businesses are in H&R.

higher covid transmission and Rufrancos et al. [2021] provide evidence that Covid-19 cases tend to spill from universities to surrounding neighborhoods.

Tourism can also encourage long-distance travel, and it has been shown that travelers contribute to the diffusion of infections. Mangrum and Niekamp [2020] provide evidence that college students returning from spring break trips accelerated the local spread of Covid-19 and the related mortality. Burlig et al. [2021] analyzes how the length of the travel ban matters for subsequent Covid-19 outcomes and shows an empirical association between migrants traveling home and the following number of cases in their home area. Finally, people attending large events with little protective behavior such as Capitol Riot (Dave et al. [2021b]) or Sturgis Motorcylce Rally (Dave et al. [2021a]) bring the disease when traveling back to their home counties.

As the literature shows, long-distance travel, and large gatherings are associated with the diffusion of Covid-19. As the opening of tourism can encourage both travels and gatherings, it is potentially highly relevant for the viral spread.

Our main contribution is the use of a unique quasi-experiment to identify the causal effect of tourism on mobility and infections. In particular, we leverage an interaction of two complementary policies which produced a large shock to touristic movement: (1) a limited-time policy that allowed opening of hotels and ski facilities in Poland during the Covid-19 pandemic and (2) the availability of travel subsidies. In addition, our study relies on novel geolocation data from Facebook, which permit us to measure mobility at a very granular scale in time and space. Furthermore, we contribute by providing quantitative answers to policymakers who seek to understand the health and economic consequences of tourism opening. We first show that opening ski resorts and hotels increased touristic spending and contributed to long-distance travel and gatherings. Secondly, we provide evidence that reopening tourism accelerated the spread of Covid-19. Finally, we show that the costs of the policy exceeded the benefits.

In the remainder of the paper, we first explain what data we use throughout the paper. The following section (3) discusses empirical methods and results regarding impact of the policy on mobility. Analogously, section 4 examines the effect of the policy on Covid-19 outcomes. The conclusion closes the paper.

2 Data

We compile a unique dataset featuring mobility patterns of Facebook users, usage of government travel subsidies, and administrative data on Covid-19 related outcomes from the Polish Ministry of Health.

2.1 Mobility

The data on mobility comes from Facebook's project Data for Good Initiative[3]. Since the start of the pandemic, it has been used for various studies mapping human mobility in countries such as the UK (Shepherd et al. [2021]), the USA (Kissler et al. [2020]), or Italy (Shtele et al. [2022], Pieroni et al. [2021], Spelta and Pagnottoni [2021], Beria and Lunkar [2021]). Data originates from Facebook users who enabled the location services on their devices. Note that manual location tagging is not required from the user. The location is captured when the Facebook or any other app using GPS is active. Users' trajectories are aggregated and anonymized to show patterns of spatial movements. We use two measures of mobility: population and collocation probabilities. Basic information on the construction of these measures is presented below, while more detailed and technical discussion has been relegated to the appendix. The reliability of the data naturally depends on the Facebook penetration of the social media market and geolocation usage. In case of Poland, Facebook is the most popular social media platform exceeding by far the competition (Hootsuite [2022], Statcounter [2022]). In early 2021, around 78% of the traffic generated by social media to other websites in Poland was from Facebook, followed by Pinterest with only 7% (Statcounter [2022]). In our data, we see about 1 900 000 users with the geolocation services turned on (around 5% of the Polish

[3]Data provided by the Facebook's Data for Good Initiative: https://dataforgood.fb.com/. We thank Alex Pompe for his help with the data

population). This number is relatively stable throughout the study period [4]. This is incomparably larger than most traditional datasets providing insights into mobility such as surveys or flight traffic data.

Nonetheless, the representativeness of such data is debatable, as specific demographics may be more likely to use social media or geolocation services. Sloan and Morgan [2015] show that Twitter users who enable geolocation have different characteristics than users who do not. Facebook users may better represent the underlying population as Gibbs et al. [2021] show that their number correlates strongly with the local census estimates in the UK. Moreover, they find no specific relationship between age, ethnicity, or poverty and Facebook usage. In the case of Poland, we see an uneven distribution in the share of the population feeding FB colocation data. As shown on the map 15b in the appendix, FB has higher penetration in western counties, which tend to be more prosperous. Overall the predictors of FB usage seem orthogonal to the location of ski resorts, and we do not expect it to change in a short time around the policy. Moreover, colocation is defined as a share of possible interactions among available users, so the measure is robust to changes in the number of users.

It needs to be acknowledged, however, that our estimates of mobility are specific to the population using Facebook in Poland. Unfortunately, FB does not share the demographic structure of its base. However, in the appendix section 7 we identify economic and demographic correlates of the consistent FB geolocation usage. Counties with high usage of geolocation tend to me more female and younger and urban, although the differences are small. It is also reassuring that FB mobility data has been shown to correlate well with other mobility sources such as geolocation from mobile operator O2 (Jeffrey et al. [2020]) or Google mobility measures (Pérez-Arnal et al. [2021]). On the other hand, Desiderio et al. [2022] use open source data (for instance train and flight traffic) to argue that FB may underestimate long-distance movements. Such bias would make estimated effects on long-distance travel conservative.

Overall, we are confident that Facebook's data provide unique and reliable insights into mobility. Below we discuss two measures used throughout the study.

Population The population at time window t and tile A is defined as the number of users who were logging mostly from the tile A during the time window t. There are three time windows per 24 hours (with breaks at 00:00, 08:00, and 16:00 UTC) and tiles are approximately 3km x 3km. Observations with less than ten users are omitted for privacy reasons. We assign tiles to counties based on their centroids. As an example, consider the map on the figure 1a. It represents the population logging on the tiles covering the Tatrzanski county - a popular tourist destination - on the afternoon (17:00 - 01:00 ETC) of the 14th of February 2021. The red dots and navy dots indicate the location of hotels and ski facilities, respectively. We use hotels and ski facilities' locations [5] to classify tiles as touristic or non-touristic. Figure 1b shows the time series of the number of users in Tatrzanski county. There is a clear uptick in the FB population after hotels' opening. See appendix for technical details, appendix table 1a and figure 13a for summary statistics of the population data.

[4] See figure 14 and the discussion on the spatial-temporal trends in FB usage in the appendix

[5] We use the name "hotel" for any accommodation facility. We find hotels and ski facilities' coordinates from the OpenStreetMaps project (OpenStreetMap contributors [2017]). The location of ski facilities were scraped from www.narty.pl and validated through own search

Figure 1: FB Population data

(a) Hotels and population in Tatrzanski county on the afternoon of the 14th February 2021

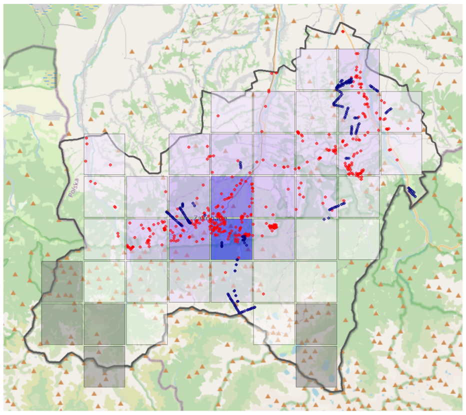

Note: The color of a tile represents the population, i.e., the number of users logging from the tile. Grey tiles correspond to no records. Red dots represent coordinates of hotels, and navy dots represent coordinates of skiing facilities. Source: OpenStreetMap and own elaboration based on Facebook data

(b) Number of FB users logging in between 17:00 and 01:00 in Tatrzanski county

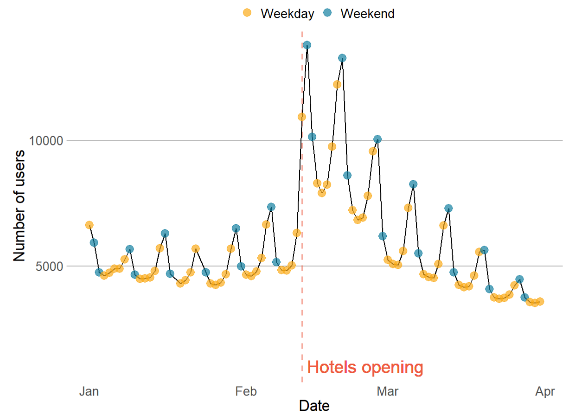

Note: The count of users in Tatrzanski county represents the sum of users from tiles with centroids in the county. Users logging in multiple tiles during the 8-hour window are assigned to the modal tile. Source: Own elaboration based on Facebook data

Colocation Colocation data aims to approximate how often users from different regions meet. Technically, it measures the probability that two randomly chosen users from county[6] i and j were within the same location, i.e., the same 0.6km $\times$ 0.6km tile [7] at a randomly chosen 5-minute interval of a given week (Wednesday to Tuesday) [8]. We do not consider collocation between users from the same county. Note that the colocation probabilities are small in magnitude because the denominator is the number of possible pairs of users from the two counties multiplied by the number of 5 minutes intervals in a week. The user's county of residence is derived from a consistent history of night-time locations. The map in figure 2 illustrates the collocation probabilities of users from Warsaw with users from other counties in the week ending on the 16th of February 2021. The size of a dot and the transparency of a link is proportional to the collocation probability between Warsaw and the given county. See appendix for technical details, appendix table 1b and figure 13b for summary statistics of the colocation data.

[6] County (*Powiat*) is an administrative unit larger than a commune. There are 380 counties in Poland

[7] Smaller tiles than in the case of population data

[8] Formally, and ignoring the week index, let X_{tir} be the number of users from region r at tile t in the 5 minute time interval i. Then let m_{rs} to be the sum of meetings between pairs of individuals from region r and s across all tiles and time intervals, that is $m_{rs} = \sum_{ti} X_{tir} X_{tis}$. The colocation probability is then the ratio of all actual meetings and all potential meetings, that is: $Pr(Colocation_{rs}) = \frac{m_{rs}}{2016 n_r n_s}$, where n_r is the number of users from region r and 2016 is the number of all five-minute intervals in a week. See Iyer et al. [2020] for more details

Figure 2: Colocation of users from Warsaw and other counties in the week ending on the 16th of February 2021

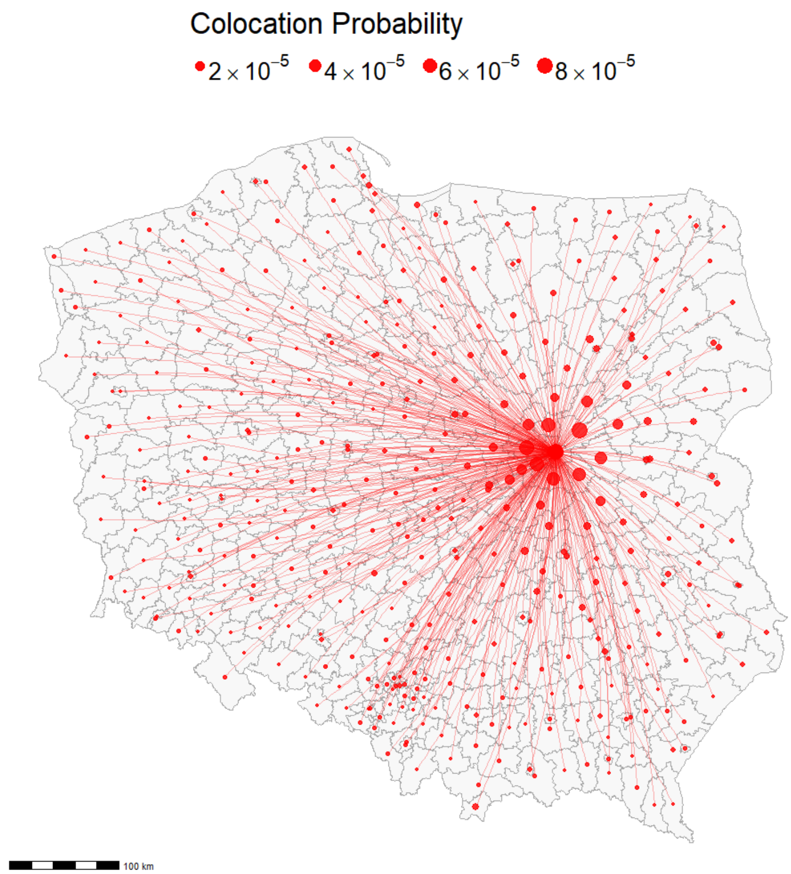

Note: Size of red dots and transparency of curved links are proportional to the colocation probabilities between users from Warsaw and given county Source: Own elaboration based on Facebook data

2.2 Spending on tourism

We approximate the spending on tourism by the usage of the funds from a governmental program subsidizing travel. The "Touristic Voucher" program (*Bon Turystyczny*) was initiated in 2020 to revive the tourism industry. Each family is entitled to one voucher per child under 18 years old. The value of the voucher is approximately \$130 [9] which can be spent on anything related to tourism, such as transportation, accommodation, or organized activities. This is a non-negligible amount given that Poles spent on average \$122 per trip in the first quarter of 2021 [10], although skiing usually requires higher spending. The data provided by the government specifies the amounts paid to businesses with vouchers. In particular, it shows total payments received by businesses in each week and each commune. Commune (*gmina*) is the lowest administrative unit in Poland, and they usually correspond to a town or a couple of villages. There are 2477 communes with a median population of 7486. The commune of the business is the commune where the headquarter is located. Data shows that the program was particularly beneficial for ski resorts. In the

[9] The value of the voucher doubles for children with disabilities
[10] Information obtained from the Polish Ministry of Sport and Tourism through the FOIA request

study period of January to April 2021, 26% of all payments made with vouchers (around \$5.5 millions) were directed to businesses located in 89 communes with ski facilities.

2.3 Health outcomes

The Polish Ministry of Health provided the data on health outcomes. It contains weekly observations at the commune level for the number of Covid-19 cases, deaths, tests taken, and people vaccinated with two doses. Our data covers May 2020-April 2021 for cases and tests, while remaining variables are available for the period January 2021-April 2021. During this period, Poland experienced its second major wave of infections. Figure 3 visualizes the the evolution of the pandemic in Poland by showing the median, 10^{th}, 25^{th}, 75^{th} and 90^{th} quantiles of weekly cases per 10 000 inhabitants across the communes. As one can see, the number of new infections varied considerably in the temporal and cross-sectional dimensions. Interestingly, the timing of the second wave coincided with the opening of hotels.

Figure 3: Quantiles of weekly cases per 10 000 inhabitants in Poland

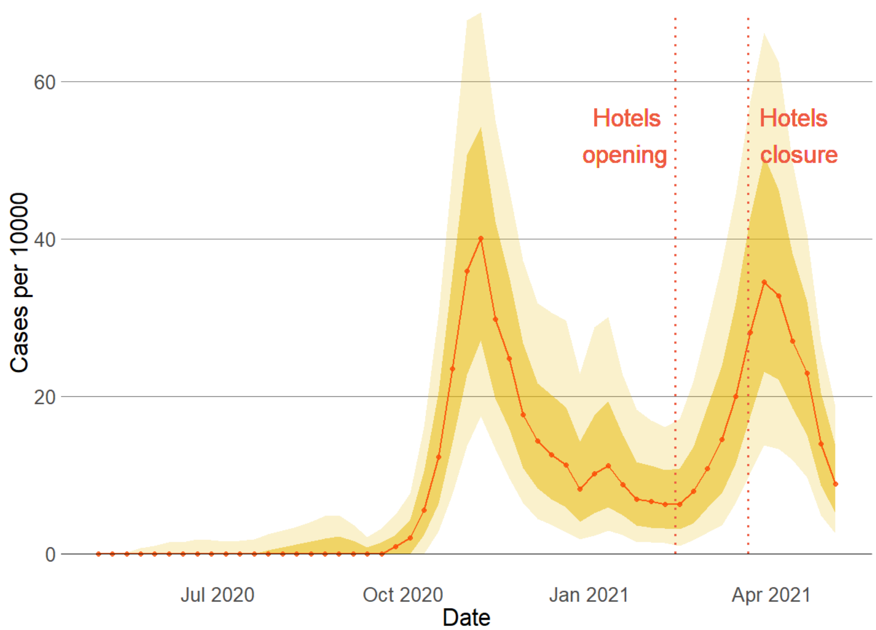

Note: The lighter shaded area corresponds to the 10^{th} and 90^{th} quantiles. The darker area corresponds to the 25^{th} and 75^{th} quantiles. The red line and points represent the median. Source: Own elaboration based on the data from the Ministry of Health

3 Mobility and spending outcomes

This section shows that the opening of hotels raised touristic spending in ski resorts and significantly increased mobility, especially at long distances. A sharp influx to touristic areas raised the frequency of meetings between inhabitants of touristic and non-touristic counties.

3.1 Empirical framework: the impact of the policy on touristic spending

We first investigate whether people responded to the policy by increasing their consumption of touristic services in the areas with ski resorts. We conduct an event study comparing spending from the "Touristic

Voucher" program in touristic vs. non-touristic areas. In particular, we divide all communes by the number of hotel beds per 100 inhabitants and the presence of ski facilities [11]. This results in a set containing three categories aiming to approximate touristic appeal: communes with fewer than four hotel beds per 100 (non-touristic), communes with more than four hotel beds but no ski resort (touristic), and communes with more than four hotel beds a ski resort (very touristic). We chose the number four as it roughly corresponds to the 90th percentile of the distribution of hotel beds per 100 inhabitants. Note that hotels and ski facilities could not operate before February 12th, so the spending trends should be parallel across categories. Then we estimate the following regression:

$$Spending_{kw} = \sum_{c \in C} \sum_{\substack{W \in \{\{01/03:01/31\}, \\ \{02/14:05/02\}\}}} Tourism_c^k I(w = W)\beta_c^W + \lambda_k + \gamma_w + \epsilon_{kw} \tag{1}$$

The event study in equation 1 analyzes the change in payments from the "Touristic Voucher" program in communes of category $c \in C$ [12] in week w compared to an analogous change in communes with fewer than 4 hotel beds per 100. The baseline period is the week ending on February 7th, which is excluded from time dummies. The outcome variable is the total value (in USD) of the vouchers spent in businesses located in commune k in week w. The dummy $Tourism_c^k$ takes value 1 if the commune k belongs to the category c. The parameter of interest is β_c^W which measures the impact of the policy on spending in areas of type c in week W compared to non-touristic areas. We expect the coefficients to be 0 before the policy and higher in areas with ski resorts after the policy as these are more appealing during winter. We allow for the commune λ_l and week γ_w fixed effects. We cluster the standard errors at the commune level.

3.2 Hotels opening led to higher consumption of touristic services in ski resorts

The opening of the tourism industry increased touristic spending, especially in the ski resort areas. Figure 4 displays β_c^W coefficients. Note β_c^W is 0 before the policy date (dashed line), confirming that trends were parallel before the reopening. However, after February 12th, we see a dramatic rise in the payments received in communes with ski facilities (right panel). The increase is smaller in touristic communes without ski resorts (right panel). The effect of the policy in ski resorts peaks in the second week and reverts to 0 as hotels close again at the end of March. Overall, the interaction of subsidies and opening of the hotels increased touristic consumption in ski resorts.

[11] Data on hotel beds come from the Polish Statistical Agency for the year 2019. Data for ski facilities come from scraping the database of the website narty.pl. See figure 17 in the appendix for the spatial distribution of ski resorts

[12] Where C contains the two non-exluded categories $C = \{$*communes with more than 4 hotel beds but no ski resort*, *communes with more than 4 hotel beds and a ski resort*$\}$

Figure 4: Event study: Hotels and ski resorts reopening and touristic spending

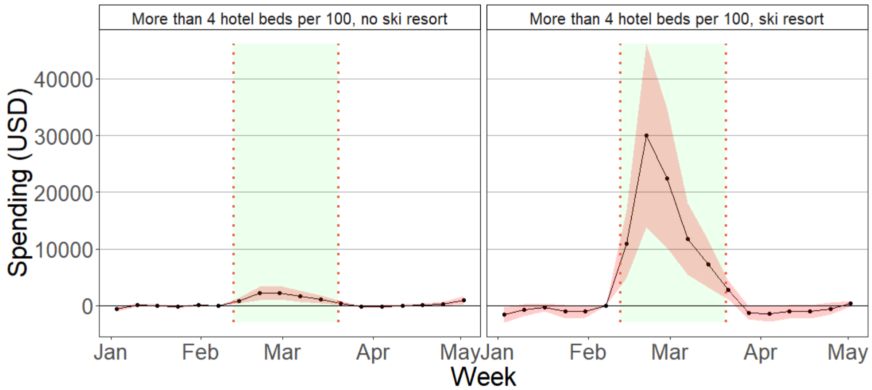

Note: Lines and points correspond to the estimates of β_c^W from equation 1. The excluded category is *Communes with fewer than 4 hotel beds per 100 inhabitants*, and the excluded date is the 7th of February. The panel on the left represents the estimates for touristic communes without ski-resorts, the panel on the right shows estimates for the touristic communes with ski-resorts. Red shaded area plots 95% confidence bands, which allows for clustering at the commune level. The green rectangle represents the time when the hotels and ski resorts were open. Source: Own elaboration based on administrative data

Some caution, however, is required in interpreting this result. Firstly, estimates present the lower bound on total spending as they account only for the payments with the vouchers. While we cannot track the expenditures from other sources, the aggregate national spending on travel in the first quarter of 2021 was 50 times higher than the amount from the vouchers only [13]. Hence, the opening likely incentivized spending also among people not using the subsidy. Nonetheless, other groups may experience different treatment effects as they differ from the voucher holders on dimensions such as age, income, or frequency of skiing, and their travels are not subsidized.

Secondly, it's worth considering if the results are robust to weather variations related to climate change. For example, warming temperatures may decrease the snow coverage and make skiing less appealing. Nevertheless, skiing resorts are popular tourist destinations throughout the year, and larger ski lifts serve hikers in the warmer months . Hence, we believe that climate variation would not fundamentally alter our results.

Thirdly, our findings do not allow us to assess what would be the impact of tourism on spending in the pre-Covid-19 period. On the one hand, domestic expenditures could be lower due to the lack of travel subsidies and the ease of border crossing before the pandemic. On the other hand, they could be higher as potential tourists did not experience adverse income shocks or did not face a risk of infection. Therefore, the pre-Covid treatment effect could vary in either direction.

While the interaction of vouchers and hotels opening makes it difficult to generalize the effect on spending, it also provides a unique setting to investigate the impact of tourism on infections. As both policies are complementary, they are likely to produce a substantial shock in the touristic movement that we can leverage to investigate its effect on public health.

3.3 Empirical framework: the impact of the policy on population movements

Informed by the increase in touristic spending, we may expect a large movement of tourists after the policy. A spike in the national railway and the passanger cars' traffic provides suggestive evidence for such movements (see figure 12b in the appendix). To investigate this formally, we use differential tourist accommodation capacity and proximity to ski resorts to conduct an event study evaluating whether the reopening of hotels increased the inflows to touristic locations. In particular, we analyze whether the number of users on

[13]The aggregate spending from vouchers in Q1 2021 was $17,359,888, while aggregate national spending on travels in the same period was $877,385,506 according to the Ministry of Tourism

tiles with many hotels increased more after the policy's implementation compared to tiles with no hotels. Furthermore, we stratify the analysis by whether the tiles are close to ski facilities. We hypothesize that the policy induced a large influx of tourists into ski resorts, noting that places with many hotels attract more tourists due to their grater capacity. Moreover, stratifying the analysis by the proximity to ski resorts partly alleviates the concern that the number of hotels proxies high urbanization as ski resorts are usually located within small towns.

To implement this strategy, we locate any accommodation and ski facilities and assign them to tiles. We then calculate the number of hotels in each tile. We bin tiles into five categories: 0 hotels, 1 hotel, 2 to 9 hotels, 10-19 hotels, and 20 or more hotels. Next, we define tiles as *in proximity to a ski resort* if they are within 25km of the closet ski lift in the mountains. Then we estimate the following regression separately for the tiles in proximity and not in proximity to ski resorts:

$$log(population)_{jtp} = \sum_{\substack{h \in \{\{1\},\{2:9\}\\ \{10:19\},\{20+\}\}}} \sum_{\substack{T \in \{\{01/06:02/03\},\\ \{02/05:03/31\}\}}} Hotels_h^j I(t=T)\beta_h^T + \lambda_{jp} + \sum_{dw \in DW} \alpha_{i(j)}^{dw} + \gamma_{tp} + \epsilon_{jtp} \tag{2}$$

The event study in equation 2 analyzes the change in the population on tiles with h hotels at date t compared to an analogous change in tiles with 0 hotels. The baseline period is the 4th of February, which is excluded from time dummies. The outcome variable is the natural logarithm of the population at tile j, date t, and time window p. The dummy $Hotels_h^j$ takes value 1 if the tile j has a number of hotels in the bin h. The parameter of interest is β_h^T, and we expect it to be 0 before the policy (12th of February), and positive after the policy. Moreover, β_h^T should increase with the number of hotels. The increase should be considerably larger in proximity to ski resorts if the movements are tourism-related. We allow for the tile×time-window fixed effects λ_{jp}, weekday×county fixed effects $\alpha_{i(j)}^{dw}$, and date×time-window fixed effects γ_{tp}. We cluster the standard errors at the county level.

3.4 Hotels opening led to increase in mobility in touristic areas

The reopening of hotels increased the population present at ski resorts considerably. Figure 5 displays β_h^T coefficients. We see that β_h^T is 0 before the policy date (dashed line), confirming that trends were parallel before the reopening in all types of tiles. However, after the 12th of February, we see high growth in the number of users present in tiles with hotels, especially in the proximity to ski resorts (left panel). The growth is also higher for tiles with more hotels. For tiles with more than 20 hotels and in proximity to ski resorts, we see about a 50% increase in the population during weekends and a 30% increase during weekdays. The effects subside with time as Poland enters its second wave of the pandemic. While there is a significant increase in population at tiles with many hotels and not in proximity to ski resorts (right bottom panel), its magnitude is small. We conclude that there was a large influx of tourists to ski resorts after the opening.

Figure 5: Event study: Hotels and ski resorts reopening and Facebook users

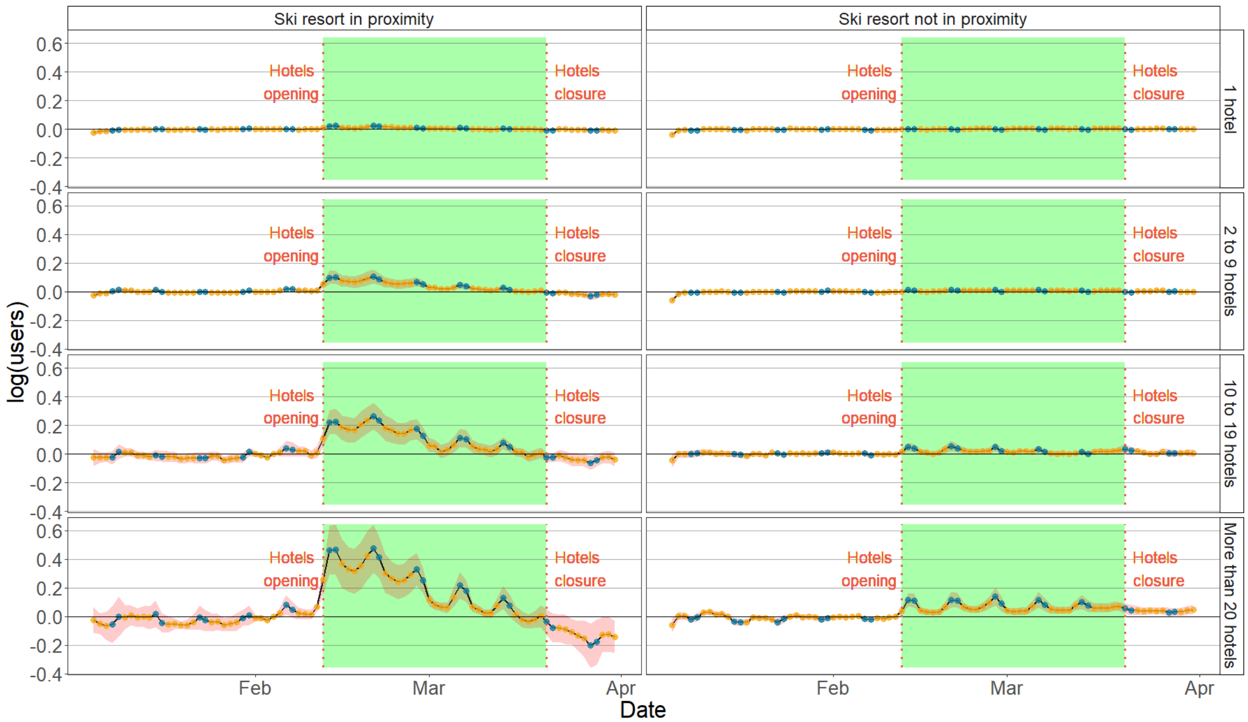

Note: Lines and points correspond to the estimates of β_h^T from equation 2. The excluded category is *tiles with 0 hotels*, and the excluded date is the 4th of February. The panel on the left represents the estimates for tiles in the proximity to ski resorts, the panel on the right shows estimates for the remainder of the tiles. Estimates of β_h^T for each bin h are plotted separately, starting with the lowest bin h at the top and the highest bin h at the bottom. Blue points correspond to weekends and yellow to weekdays. Red shaded area plots 95% confidence bands, which allows for clustering at the county level. The green rectangle represents the time when the hotels and ski resorts were open. Source: Own elaboration based on Facebook data

3.5 Empirical framework: the impact of the policy on the frequency of meetings

In this section, we investigate whether the policy affected the frequency of meetings between users from different counties. Such meetings are essential from an epidemiological perspective because they can transform local outbreaks into a national wave. Our hypothesis is that the reopening policy made population flows at long distances and flows from non-touristic to touristic counties more likely. To test this, we perform two analyses. The first analysis investigates whether the frequency of long-distance meetings increased relative to short-distance meetings after the policy was enacted. We classify each link into five distance bins based on the distance between centroids of the counties. The bins are $d \in D = 0-100km, 100-200km, 200-300km, 300-400km, 400+km$. Next, we regress the log of colocation probabilities on the interaction of the week dummies with the distance bins:

$$log(P(collocation))_{klw} = \sum_{\substack{W \in \{\{01/05:02/02\},\\ \{02/16:04/13\}\}}} \sum_{d \in \underline{D}} Distance_{kl}^{d} I(w=W)\beta_d^W + \phi_{kl} + \chi_w + \upsilon_{klw} \tag{3}$$

Where $log(P(collocation))_{klw}$ is the log of the probability of collocation between users from county k and l in the week w. A dummy $Distance_{kl}^{d}$ is equal to 1 if the distance between counties l and k is in the bin d.

The bin with the shortest distance is excluded as a reference. The dummy $I(w = W)$ is equal to one if the week of the observation corresponds to the week W. The excluded week is the last week before the opening, that is, the week of 02/09. We allow for link ϕ_{kl} and week fixed effects χ_w, and we cluster the standard errors at the link level. The parameter of interest is β_d^W which is a difference-in-differences estimator: the first difference measures the percentage change in the collocation probabilities for counties at a distance d in the week w compared to the week of 02/09. The second difference takes this change and compares it to an analogous change for counties at a distance within 0-100km. We expect β_d^W to be positive for weeks when the policy was in place as people started to travel long distances.

The second analysis tests the hypothesis that meetings between non-touristic and touristic counties increased after the policy was implemented. We rely on the assumption that counties with large skiing and accommodation capacities have greater touristic appeal (see the derivations of the theoretical impact of policy on colocation in the appendix). Hence, we define the exposure to tourists by using the total length of skiing trails and the number of hotel beds in the county[14]. In particular, we first classified counties as below ($few\ hotels$) or above ($many\ hotels$), the third quartile of the distribution of hotel beds. Second, we classify counties with 0 skiing trails as 0 $trials$. Finally, for counties with some skiing trails, we divide them by whether they are below ($few\ trails$) or above (($many\ trails$) the third quartile of the distribution of the total length of skiing trails [15]. In this way, we obtain five possible exposure statuses $es \in ES$={*0 trails & few hotels beds, 0 trails & many hotel beds, Few trails & few hotel beds Few trails & many hotel beds, Many trails & Many hotel beds*} [16]. See figure 18 for the spatial distribution of the exposures.

We expect the probability of collocation between inhabitants of touristic and non-touristic regions to have increased after the policy's implementation. The most substantial effect should exist for pair of counties with and without ski resorts. To measure such effect, we conduct our analysis on the link level (links between counties) by running the following regression:

$$\begin{aligned} log(P(collocation))_{klw} = &\sum_{\substack{W \in \{\{01/05:02/02\},\\ \{02/16:04/13\}\}}} \sum_{s \in \underline{ES}} \sum_{q \in \underline{ES}} Exposure_k^s \times Exposure_l^q I(w = W)\beta_{sq}^W \\ &+ \delta_{kl} + \gamma_w + \epsilon_{klw} \end{aligned} \tag{4}$$

A dummy $Exposure_k^s$ is equal to 1 if the county k belongs to the exposure category s and 0 otherwise. The excluded combination of classes is the one between counties which both belong to *0 trails & few hotels beds*. Analogously to equation 3, we exclude the week of 02/09, and we allow for link δ_{kl} and week fixed effects γ_w, and we cluster the standard errors at the link level. The parameter of interest β_{qs}^W estimates the percentage change in the collocation probabilities between users from counties of types s and q in the week w compared to the week of 02/09 and relative to an analogous change for users from two different counties both belonging to the type *0 trails & few hotels beds.* Note that the links are undirected, hence we use β_{qs}^W independently of which county is in q and which in s. We expect β_{qs}^W to be positive after policy implementation for pairs s and q such that one has ski trails and hotels while the other does not. Moreover, the effect should grow in the difference between the counties touristic appeals and hence should see the most prominent effect for pair *0 trails & few hotels beds* and *Many trails & Many hotel beds.*

3.6 Policy increased likelihood of long distance meetings and meetings between locals and tourists

The frequency of long-distance meetings and meetings between non-touristic and touristic counties increased after the opening. Figure 6 shows the parameter of interest from the equation 3. We see a clear increase in colocations at long distances compared to counties within 100km after the policy was enacted. Moreover,

[14] Data from the Polish Statistical Agency for the year 2019
[15] Among counties with any skiing trails
[16] All counties with long skiing trails have many hotels beds, hence exposure *Few hotel beds & Many trails* is missing

the increase was greater for distances above 200km, which is consistent with tourists going further as they can stay for the night in a hotel. The parameter of interest decreases after the initial surge, which may be related to the rising number of Covid cases in late March. Furthermore, we see a spike in the first week of April, which likely corresponds to Easter festivities.

Figure 6: Event study: Hotels and ski resorts reopening and long-distance colocation

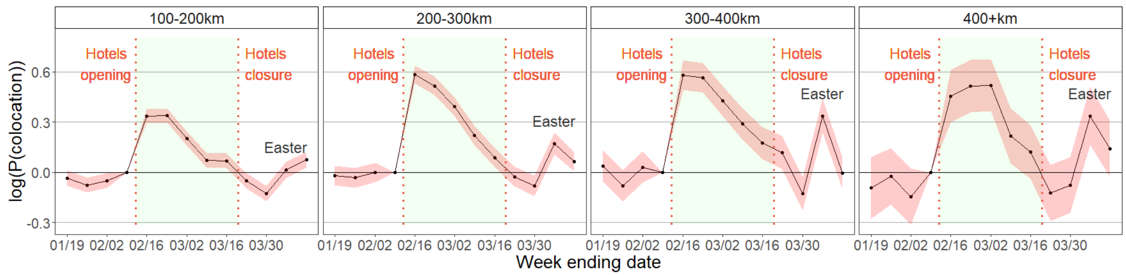

Note: Lines and points correspond to the estimates of β_d^W from equation 3. The excluded category is *Distance < 100km* and the excluded week is 02/09. Each panel represents estimates of β_d^W for a different distance bin b starting with the lowest distance on the left. Red shaded area plots 95% confidence bands, which allows for clustering at the link level. The green rectangle represents the time when the hotels and ski resorts were open. Additionally, an annotation is added to mark the week of Easter. Source: Own elaboration based on Facebook data

Turning our attention to equation 4, figure 7 shows an increase in colocation between tourists and locals. Each panel in figure 7 corresponds to estimates of parameters β_{qs}^W for different s and q. The row label represents category s, and the column label represents category q. For example, the top-left panel represents the change in the colocation probabilities between counties with *0 trails & few hotels beds* and with *0 trails & many hotels beds*. As expected, we do not see significant changes in the frequency of meetings between pairs of counties that are either both non-touristic (that is, have 0 or few trails) or both touristic. While the null effect among non-touristic counties is accurately estimated, we obtain noisy estimates for links between counties that both have trails. This is due to a lower number of touristic counties and, hence, fewer connections among them. Importantly, we see a significant increase in the probability of meetings between counties with 0 trials and counties with many trials immediately after the opening. The magnitude of this increase is around 50% in the week following the the opening and stays positive for three weeks.

Figure 7: Event study: Hotels and ski resorts reopening and touristic colocation

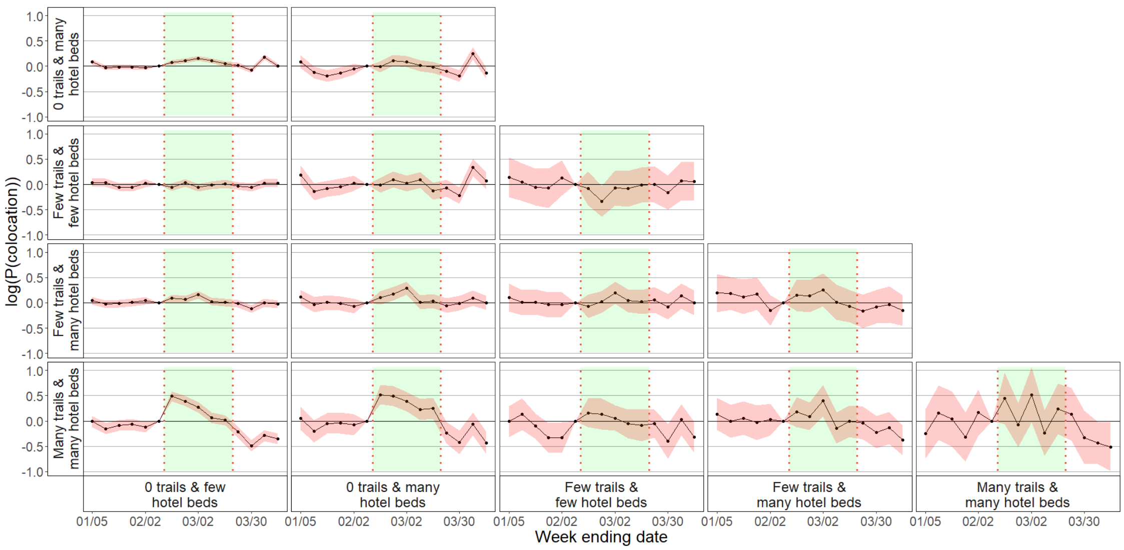

Note: Lines and points correspond to the estimates of β^W_{sk} from equation 4. The excluded category is one with both counties belonging to *0 trails & few hotels beds* and the excluded week is 02/09. Each panel represents estimates of β^W_{sk} for a different pair of s and k types. Note that the ordering of types does not matter because links are symmetric. The types are described in the strips on the left and on the bottom. For example, the bottom left panel represents β^W_{sk} where one county belongs to *0 trails & few hotels beds* and the other to *many trails & many hotels beds.* Red shaded area plots 95% confidence bands, which allows for clustering at the link level. The green rectangle represents the time when the hotels and ski resorts were open. Additionally, we add an annotation to mark the week of easter. Source: Own elaboration based on Facebook data

We conclude that the policy increased the frequency of long-distance meetings and that of meetings related to tourism. As such, it could have a significant impact on Covid-19 outcomes.

4 Covid-19 outcomes

4.1 Conceptual framework

As the policy encourages gatherings and travel, it should impact the number of infections by increasing the number of contacts between individuals. However, towns (communes) should be affected deferentially depending on their participation in the tourism industry. We exploit this exposure heterogeneity when designing our identification strategy.

Conceptually, we can distinguish between three levels of the exposure to the policy. First, communes hosting ski facilities (type 1) are highly exposed. They experience a large influx of tourists as well as an increase in local interactions due to amplified economic activity. We may expect new cases arising among locals who come in contact with a higher number of individuals. Second, communes sending tourists to ski resorts (type 2) are also directly affected by the policy. Tourists come into contact with both locals and other tourists while staying at the resorts, and consequently they are at a higher risk of getting infected. Third, there are communes that did not send tourists to ski resorts after hotels opened (type 3). They are the least exposed to the policy as their contact patterns have not changed.

While the contact patterns at type 3 communes have not been affected, these communes can still experience treatment effect in terms of new infections. Policy-induced cases in communes of type 1 and 2 can produce secondary infections which spread through the existing networks (Chang et al. [2021],Kuchler et al. [2022],Fritz and Kauermann [2021]). They can flow to communes not sending tourists to ski resorts through the connections pre-existing the policy. Note that these additional secondary cases would not have happened in the absence of the policy and hence they are part of the treatment effect.

Although every commune can be potentially affected by the policy, we may expect that the timing of the treatment effect differs by type (Shtele et al. [2022], Thomas et al. [2020]). In particular, policy induced infections should first appear in communes hosting ski resorts or directly sending tourists, and only later in communes not sending tourists.

Motivated by this reasoning, our identification strategy leverages differential exposure to the policy by type of the commune and time after the opening. In particular, it relies on comparing the dynamic of infections in communes directly exposed to the policy (type 1 and type 2) to the analogous dynamic in the communes not sending tourists to ski resorts.

4.2 Empirical framework

Translating this conceptual framework to the empirical framework poses three main challenges. First, we need to identify communes sending tourists. Second, we need to adjust for differences in the pre-policy pandemic outcomes across communes. Third, we need to find a valid counterfactual in the realm of a nation-wide policy.

4.2.1 Tourists origin communes

We use colocation data to identify communes sending tourists to ski resorts. In particular, we sum all meetings between individuals from county i and any county containing ski resort [17] during the first three weeks of the policy. Mathematically, the strength of connections, which we call *exposure* is:

$$exposure_i = \frac{\sum_{w \in \{w*, w*+1, w*+2\}} \sum_{j \in ski\_resort} Meetings_{ijw}}{N_i}$$

where N_i is the number of FB users in county i and ski_resort is the set of counties with ski resorts. Note that the connection measure is at a lower geographic granularity (county) than the outcome (commune). Hence, we assume that all communes in a county have the same $exposure$. We divide communes into 5 categories which correspond to types. Communes with ski resorts are classified as *Ski resort commune*. Communes without ski facilities but in the counties with ski resorts are classified as *Ski resort county* [18]. The rest of the communes are classified according to the tertiles of the *exposure* of their county (similarly to Dave et al. [2021b]). We assume that the counties in the first tertile did not send any tourists because they had fewest interactions with inhabitants of ski resorts during the policy. We further assume that counties in the third tertile sent more tourists than counties in the second tertile. The map on figure 8 shows the spatial distribution of the treatments [19].

[17] We exclude ski facilities not located in the mountains

[18] This distinction is necessary because we do not know connections within the county

[19] Note that the connections are at the county levels, but the outcomes are at the commune level. Hence, communes in the same counties will have the same exposure

Figure 8: Exposure to ski resorts

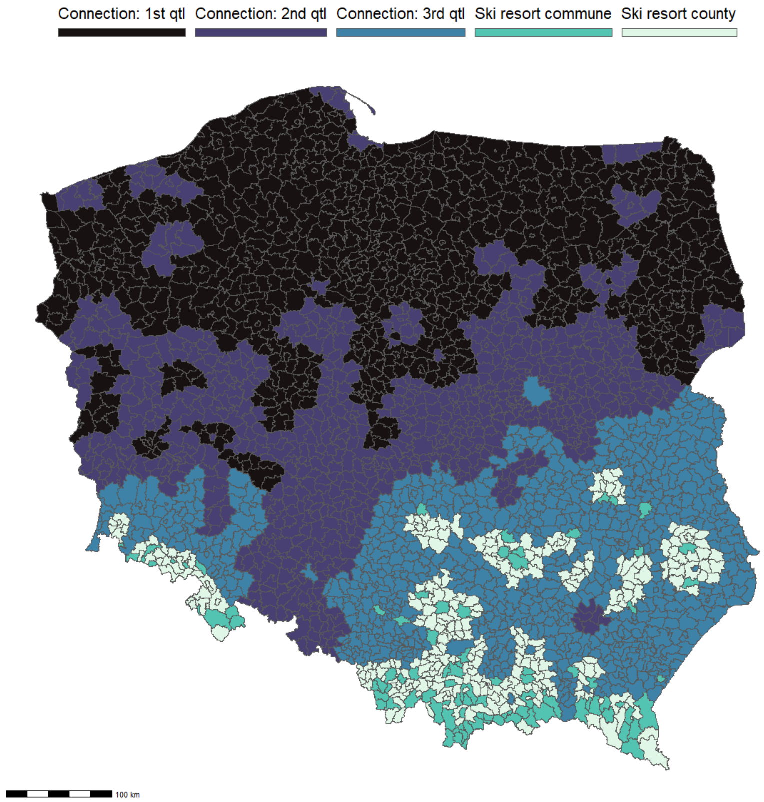

Note: Each color corresponds to a different exposure category. Communes within the same county belong to the same exposure unless the county contains a ski resort. Communes containing ski facilities are in the category *Ski resort commune* and communes without ski facilities but in counties containing ski resort communes are in the category *Ski resort county*. Source: Own elaboration

For the preliminary empirical evidence on the impact of the policy, we compare new cases in communes with different treatments. Figure 9a plots the average number of new cases per 10 000 inhabitants in communes by their exposure to the policy. Firstly, one notes that the pre-policy infection trends differ between the treatment arms. Secondly, we do not see a significant difference in the number of new infections right after the enacted policy (first dotted line). However, such a simple comparison may be insufficient to uncover the impact of the policy. The diffusion is affected not only by the number of contacts but also by the number of previous infections and susceptible individuals. Different treatments seem indeed to be at a different stage of the pandemic before the policy, as evidenced in figure 9a. This by itself is enough to render a simple comparison unreliable.

4.2.2 Balancing communes by pre-policy outcomes

To analyze the policy, we need to compare locations at a similar pandemic stage before the policy. As SIRD has markovian properties, the current pandemic situation should depend only on the indices in the previous period. The variables representing number of infections, share of susceptibles and number of infections among connected units in the current period act as sufficient statistics for the outcome in the next period. Consequently, the communes with similar pre-pandemic outcomes and characteristics should evolve in the same way in the absence of the policy. Hence, our analysis conditions on a set of pre-policy variables $F_{k,w*-1} = \{X_{k,w*-1}, Z_k\}$. In particular, $X_{k,w*-1}$ represents pandemic outcomes in the last period before the policy. It includes the number of cases, deaths, and tests per 10 000, their growth rates, and their cumulative numbers since the beginning of the pandemic to proxy for share susceptible[20]. It also contains the weighted sum of cases among neighbors where the weights correspond to the number of commuters per capita. Moreover, $X_{k,w*-1}$ includes the squares and interactions of all these variables. The controls Z_k represent observed location-specific characteristics which may influence the diffusion. In particular, Z_k includes population, population density, the type of the commune (urban or rural), unemployment rate at the end of 2020, the number of sport objects per capita (swimming pools, stadia, courts), and the number of theatres and cinemas per capita. The remaining analysis relies on the assumption that in the absence of the policy and conditional on these variables, the potential outcomes would evolve in the same way across all the types:

Assumption 1 *-Ignorability* : $Y(0)_{k,w\geq w*} \perp\!\!\!\perp Treatment_k^{tp} | X_{k,w*-1}, Z_k$
for $tp \in \{ski\ resort, sending\ tourists, not\ sending\ tourists\}$

Figure 9b plots the average number of new cases by type in the sample matched on the conditioning variables. We notice a considerable improvement in the similarity of the pre-policy outcomes, suggesting that the matching was successful. However, the outcomes diverge in the weeks after the opening. In particular, the number of cases is higher in the first weeks after the opening in locations with ski resorts and communes strongly connected to ski resorts.

Figure 9: Covid cases in ski resorts

(a) Covid cases: full sample

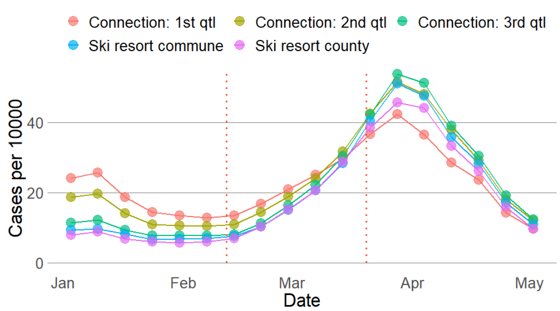

Note: The averages were calculated on the full sample. Dotted lines represent the opening and closure of hotels. The date corresponds to the last day of the week. Source: Own elaboration based on the Ministry of Health Data

(b) Covid cases: matched sample

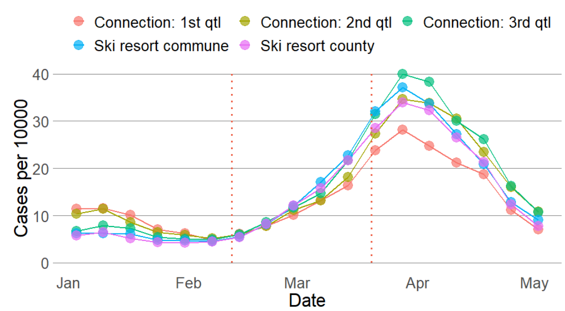

Note: The averages were calculated on the matched sample. Each treated unit was matched with one unit from the first tertile of connection. Units were matched by the distance in the propensity scores computed based on the conditioning variables. Dotted lines represent the opening and closure of hotels. The date correspond to the last day of the week. Source: Own elaboration based on the Ministry of Health Data

4.2.3 Identifying control communes

The policy is nationwide, however the exposure to the policy differed between the communes. We argue that the least exposed communes can act as controls to identify the *on impact* effect of the policy. Since the level of interactions did not change in communes not sending tourists, their infection dynamic was not affected at the onset of hotels opening. Consequently, they can serve as controls in the initial period of the new policy. Hence, we are able to identify *on impact* treatment effect in communes with ski resorts and analogously in communes sending tourists. The identification is more challenging in future periods. As the time passes, infections induced by the policy flow through the existing network of connections. Due to these spillovers, the treatment effect is no longer null in communes not sending tourists in later periods. In particular, they

[20] Since the beginning of 2021 in case of deaths, as the number of deaths is only available for 2021

are affected by higher level of infections in communes that they are connected to. Hence the comparison is biased in later periods and the bias is larger if communes are well connected and hence experience large spillovers. However, the treatment effect in communes not sending tourists is presumably non-negative as the policy should increase number of new cases. Hence, we can identify lower bounds on the treatment effect in the remaining communes.

4.2.4 Estimation

To estimate the effects, we turn to the unconfoundedness approach suggested by Callaway and Li [2021]. It has been shown that fixed effects estimation is unlikely to produce reliable results if outcomes are generated by a non-linear model (Gauthier [2021], Callaway and Li [2021]), and the infections are likely a product of such non-linear model[21] (Keeling and Eames [2005], Brauer [2017], Caccavo [2020]). This issue occurs because the treated units can be at a different pandemic stage than the control units. For example, suppose that the treated units experienced their first case earlier than the control units. Then, their outcomes will not evolve in parallel to the outcomes in the control areas, even in the absence of any policy. Figure 10 illustrates this problem. It plots results of a simple simulation of the Susceptible-Infected-Recovered-Deceased model in two areas which are identical except for the timing of their first case. The area represented with the red curve started the pandemic earlier. Dotted lines show linear trends in the number of cases between two points in time. Despite identical parameters, the trends are not parallel because these areas are at different pandemic stages. Hence, differencing trends would introduce a bias rather than remove it.

Figure 10: Non linearities implied by SIRD model

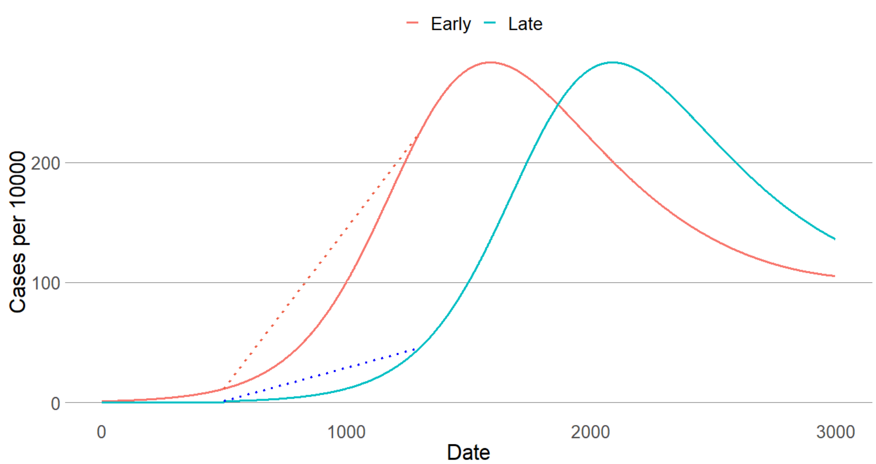

Note: figure plots the results of a simulation of SIRD model in two areas. The parameters are identical in both cases, however the timing of the first case is different. The area represented by red curve had the first case earlier than the area represented by the blue curve. Dotted lines measure linear trends in cases between two points in time. Despite identical parameters, areas experience non-parallel trends

The unconfoundedness approach alleviates the above issue in two ways. Firstly, it does not rely on fixed effects. Secondly, it ensures that the control units are at a similar pandemic stage before the policy. This is achieved by conditioning on the pre-treatment covariates related to the pandemic. Thus, this approach is compatible with a case in which pandemic-related parameters vary over time and with commune-specific characteristics. It is, however, not compatible with a general unobserved heterogeneity in parameters by

[21]We perform traditional event studies on the full (figure 16a) and matched samples (figure 16b), see the results in the appendix

location. Hence, the estimation presented below is valid under the assumption that the parameters change across locations only due to the variation in the conditioning controls $F_{k,w*-1}$.

Intuitively, the unconfoundedness approach computes a weighted difference of outcomes in treated communes versus control communes which have similar pre-treatment characteristics to the treated units. The larger the similarity is, the higher is the weight for the control unit. In particular, denote the pre-treatment characteristics as $F_{k,w*-1} = \{X_{k,w*-1}, Z_k\}$ and $C_{w,k}$ as the number of new cases in week w and commune k. Following the notation from Callaway and Li [2021], we estimate:

$$ATT_w^c = E[w(treatment_k^c, F_{k,w*-1})(C_{w,k} - m_{0,w}^C(F_{k,w*-1}) \tag{5}$$

where ATT_w^c corresponds to the average effect on the treated by treatment c, and the weights correspond to:

$$w(treatment_k^c, F_{k,w*-1}) = \frac{treatment_k^c}{E[treatment_k^c]} - \frac{\frac{p(k,F_{w*-1})}{1-p(k,F_{w*-1})}(1 - treatment_k^c)}{\left[\frac{p(F_{k,w*-1})}{1-p(F_{k,w*-1})}(1 - treatment_k^c)\right]}$$

Finally, the untreated potential outcomes correspond to

$$m_{0,w}^C(F_{w*-1}) = E[C_w | F_{k,w*-1}, treatment_k^c = 0]$$

.

The untreated potential outcome of a treated unit k in a week w comes from outcomes of untreated units similar to k in week w. Their expectation is unbiased for the untreated potential outcome under the assumption 1. Concretely, we implement this method by estimating propensity scores $p(F_{k,w*-1})$ with logit and outcome regression $m_{0,w}^C(F_{k,w*-1})$ with OLS. The method is double robust because it is robust to the misspecification in either propensity scores or conterfactual regression of potential outcomes. The method relies on the assumption that the pandemic trajectory would be similar in the treatment and control groups after conditioning on the $F_{k,w*-1}$. We argue for the validity of this assumption in the section 4.2.5. Note that we estimate four effects: one for each treatment versus the control group of communes in the first tertile of *exposure*. Hence, we estimate the effect in four samples where each sample contains units from the control and one of four treatments. Figure 11 plots the results of the estimation together with 95% confidence bands calculated with multiplier bootstrap.

Figure 11: Event Study: Covid-19 and tourism opening

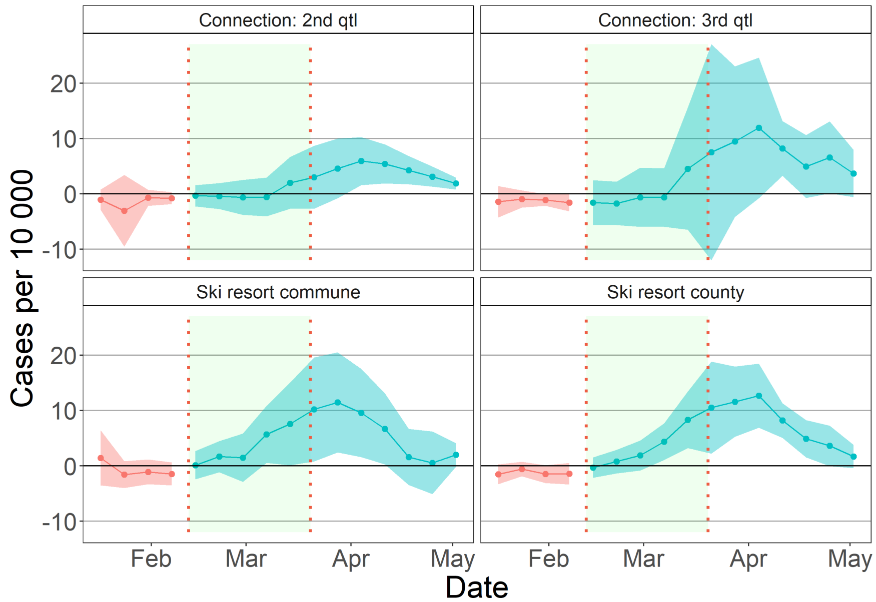

Note: The estimates come from the unconfoundedness approach by Callaway and Li [2021]. Red points correspond to the estimates for pre-treatment periods. The reference point in a pre-treatment period t is the previous period $t-1$. Blue points correspond to the estimates for post-treatment periods. The post-treatment periods' propensity scores and outcome regression are based on the last period before the policy $w*-1$. Control units are communes in the first tertile of *exposure*. The shaded area represents simultaneous 95% confidence bands with clustering at the commune level. The date corresponds to the last day of the week.

The results are consistent with the hypothesis that the policy precipitated the arrival of the infections' wave, and that tourists brought the disease back to their home counties. Namely, communes in counties with ski resorts experienced a higher number of infections after the opening. They have about five additional infections per 10 000 in the third week after the opening comparing to the control. The effect is statistically and economically significant as it represents a 70% increase with respect to the pre-policy average in these communes. We interpret this as the *on impact* treatment effect. It is not obvious why communes without ski facilities but in counties with ski resorts experienced impact equivalent to ski resort communes. There may be many interactions between the communes with ski resorts and without ski resorts in the same county. For instance, people may be going skiing or to work in nearby ski resorts. Figure 19 in the appendix shows that commuting to ski resorts communes is stronger than across no ski-resort communes, but magnitude of commutes are not large.

The early increase in cases is absent in communes in the second and the third tertile of *exposure*. However, communes with medium and strong connections see an increase in infections compared to control starting in the fourth week after the opening. This could be the result of secondary infections from tourists bringing the disease back home and spillovers from other communes. Moreover, we see a monotonicity of this effect in the strength of connection: communes in the third tertile have a higher increase in cases than communes in the second tertile. While this is consistent with the story of tourists contributing to the diffusion of the virus, these differences are not statistically significant.

The above analysis provides suggestive evidence that the opening of hotels contributed to the diffusion of COVID-19 through touristic gatherings and travels. In particular, the wave of infections arrived earlier in the counties with ski resorts. Moreover, there were more infections in counties with ski resorts and counties strongly connected to ski resorts than counties weakly connected to ski resorts. Note, however, that this exercise does not allow us to conclude whether the opening of tourism caused or did not cause the second wave of the pandemic in Poland. While we find differential trends in infections by exposure to tourism, there are no units that would not be affected by the policy at medium or long term. Hence, there is no plausible counterfactual which would allow for the evaluation what would happen in the absence of the policy at medium or long term.

4.2.5 Robustness

The method developed by Callaway and Li [2021] relies on the assumption that applying the propensity score weights and the regression to control units can predict the counterfactual outcomes for treated units. In other words, control-based predictions should replicate the potential trajectory of infections in the treated communes that would have happened in the absence of the policy. While it is impossible to test the counterfactual's performance during the time of the intervention, one can look at periods before the policy's enactment. In particular, conditional on $F_{k,w*-1}$, there should be no difference in the outcomes between treated and control communes before February 12th.

Hence, as the first check, we extend the study period until November 2020 to examine the pre-policy differences. We compare the counterfactual predictions[22] and the treated outcomes before the implementation date and we find no considerable differences as shown on the figure 21 in the appendix.

As the second check, we perform a placebo exercise where we set the treatment timing to start a month before the actual implementation date. We expect no differences between the "treated" trajectory and the control-based predictions as the actual policy has not started yet[23]. Figure 22 confirms this intuition showing no major differences between the treated and control communes after the placebo date. While there are some deviations from 0, they are small and opposite to the policy effect. Hence, we believe the method provides a reasonable counterfactual trajectory for the treated communes.

In addition to the opening of hotels and ski resorts, the government changed some other restrictions during the study period. We take steps to ensure that these additional changes are not driving our results. Firstly, simultaneously to lifting the closure of hotels, Polish authorities allowed theatres, cinemas, swimming pools, and outdoor sports venues to resume their activity. For this reason, the main specification includes the per capita number of theatres, cinemas, and sports venues [24] as conditioning variables. Consequently, treatments and control groups should be balanced with respect to the availability of these venues; hence, this should not affect their infection dynamics deferentially. Secondly, the Polish government introduced stricter measures in four regions [25]: on the 27th of February in Warminsko-Mazurskie, and then on the 15th of March in *Pomorskie*, *Mazowsze*, and *Lubuskie*. The measures included closing hotels, theatres, sports venues, malls, and remote learning for primary schools. These measures covered the entire country starting on the 20th of March. Note that none of these regions contains ski resorts. As a robustness check, we repeat our analysis excluding these four regions from the sample. The results (presented in the appendix in the figure 20) are qualitatively unchanged, albeit less precise. Finally, some nationwide changes were implemented, such as mandatory quarantine for certain international travelers and mandatory covering of the face with a medical-grade mask (as opposed to bandanas). As these measures are national, we do not expect they would affect our treatment and control groups differently.

5 Cost-Benefit Analysis

While opening the hotels and ski lifts revived the tourism industry, it also produced public health costs. We evaluate the policy by quantifying its costs and show that they are larger than the benefits.

[22] Note that the pre-treatment counterfactual predictions are based on the preceding period, while post-treatment predictions are based on the last period before the policy, which results in a longer prediction horizon

[23] The main difference with the previous check is a longer prediction horizon for the counterfactual outcomes trajectory

[24] As reported by the Central Statistical Office.

[25] There are in total 16 regions in Poland.

We assume that the primary costs stem from increased usage of healthcare resources and deaths. Consider first the hospital costs. Accounting for the probability of hospitalization from Covid-19, its length, and fees in Poland (based on Orlewska et al. [2021]), each case has an expected healthcare expenditure worth $275. Next, consider the cost of Covid related deaths. Computing the cost of lives lost requires identifying fatality likelihood and assigning a value to each life. The fatality rate of Covid-19 in the period of interest in Poland is 2.7% [26]. We use information provided in Robinson et al. [2021] to monetize the value of life. In particular, they combine the constant value per statistical life-year (VSLY) and age distribution among Covid-19 deaths to calculate the expected cost of years lost due to a Covid infection. They assign $4.47 million per life lost to Covid, which results in an expected death cost of $120 690 per infection. We chose this measure for two reasons. Firstly, Covid-19 deaths are concentrated among the elderly, and the VSLY measure accounts for their older age. Secondly, this technique gives the lowest estimate among other methods proposed by Robinson et al. [2021], and hence it helps construct a lower bound on costs. Summing up the healthcare expenditures and the value of life lost, each case is associated with an expected cost of $120,965.

The policy caused 134,886 additional infections throughout the country. We obtain this number by summing the estimated treatment effect across populations of all affected areas in the periods after the implementation. Note that given the 2.7% fatality rate, we expect 3642 deaths were due to the opening. The policy's total cost is the product of infections times their expected cost, which is $16.316 billion.

Estimating the benefits is more challenging because we only have granular data on spending from the vouchers. Nevertheless, we can identify the upper bound of policy's contribution to the GDP. According to the Polish Ministry of Sport and Tourism, touristic expenditures in the first quarter of 2021 totaled $0.877 billion. Even if all this spending stemmed from the policy, it would still be only about one twentieth of the lower bound on the cost. Hence we conclude that the policy's cost vastly surpassed its benefits.

The policy benefited only tourists coming to the ski resorts. By the revealed preferences, we conclude that they have enjoyed a positive surplus because they chose to go skiing despite the risk of infections. On the other hand, people outside ski resorts who are not tourists were worse off because they experienced negative externalities of the policy without any benefits. A priori, the results for the ski resorts' inhabitants are ambiguous. From our results it is clear that they did not benefit from the policy. It caused 13,448 new cases in communes with ski facilities which led to a cost of $1.626 billion. This number is still higher than the upper bound on the benefits. The costs would surpass policy's benefits as long as the number of produced infections exceeded 7250, or alternatively as long as there were more than 197 deaths. Note that a wider access to vaccinations and treatments could potentially make the policy beneficial. However, at the time of the opening only 1.4% of Poles were fully vaccinated.

Tourism during the pandemic is a risky behavior. It exposes tourists to the virus and subsequently it contributes to a wider diffusion of cases among the general population. This negative externality produced costs which vastly outsized the benefits of hotels opening. Even the ski resorts were worse off, despite the intent of the policy. The only beneficiaries were tourists, particularly those with children, who could enjoy subsidized travel. Both the voucher and the opening may be seen as a transfer from the general population to the families with children, which the Polish Government tends to favor.

6 Conclusion

Tourism plays a vital role in providing income for many local economies. However, while important for economic reasons, it also encourages long-distance travel and gatherings. Moreover, touristic services often require risky in-person interactions. These features make tourism a transmission vector for various infectious diseases. We hope that our analysis will provide some guidance for policymakers struggling to balance the trade-off between economic and public health goals related to the opening of the tourism industry.

In this paper, we investigated how the opening of hotels and ski facilities impacted touristic consumption, mobility and Covid-19 outcomes. The opening was followed by large movements of tourists to locations with ski facilities. Areas with many hotels in proximity to the ski trails experienced an exceptionally high influx of

[26] calculated as the ratio of deaths to cases lagged by two weeks

visitors and spending. Travels often originated from distant locations, and hence the probability of meetings between individuals residing far from each other increased after the policy. Additionally, there has been an increase in meetings between pairs of individuals such that one person lives in a touristic and one in a non-touristic location. These observations point out the strong impact of tourism opening on mobility.

Travelers have a high potential to carry the disease between distant locations. This is particularly dangerous when they also participate in gatherings. Visitors of ski resorts could only gather in their hotel rooms and on trails as the restaurants were closed. Nonetheless, there is suggestive evidence that they impacted the Covid-19 trajectory. We showed that having a ski facility in a county is correlated with an increase in infections after the policy. Moreover, counties with frequent meetings with ski resorts during the opening had more infections than counties with few such meetings.

We believe that our results can be extrapolated to other settings involving mass tourism. We think that travels and gatherings are the main factors driving additional infections related to tourism opening. As long as these two elements are present, one may expect an increase in the number of cases, although of different magnitude depending on the circumstances. While the effects may be more substantial during winter because people spend more time indoors, there is still a considerable amount of close interactions in touristic activities during other seasons. For instance, travelers still share public means of transportation, and locals engage in repeated interactions with tourists when providing them services. Hence, additional opportunities for transmission still arise.

Our study shows that engaging in touristic activity can generate negative externalities as it contributes to the spread of infections. Hence, it might be reasonable to impose some additional costs, such as post-travel quarantine, for people involved in tourism.

We note that the policy was enacted before the full distribution of Covid-19 vaccines. These could potentially mitigate the impact of tourism on Covid-19. Nonetheless, long-distance travellers have a high potential to carry novel variants to new locations. Future research could explore whether tourism activity is associated with a faster arrival of new variants.

References

Rahi Abouk and Babak Heydari. The Immediate Effect of COVID-19 Policies on Social Distancing Behavior in the United States. *medRxiv*, page 2020.04.07.20057356, April 2020. doi: 10.1101/2020.04.07.20057356. URL `https://www.medrxiv.org/content/10.1101/2020.04.07.20057356v2`.

Daron Acemoglu, Victor Chernozhukov, Iván Werning, and Michael D. Whinston. Optimal Targeted Lockdowns in a Multigroup SIR Model. *American Economic Review: Insights*, 3(4):487–502, December 2021. doi: 10.1257/aeri.20200590. URL `https://www.aeaweb.org/articles?id=10.1257/aeri.20200590`.

Jérôme Adda. Economic Activity and the Spread of Viral Diseases: Evidence from High Frequency Data *. *The Quarterly Journal of Economics*, 131(2):891–941, May 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw005. URL `https://doi.org/10.1093/qje/qjw005`.

Ahammer Alexander, Halla Martin, and Mario Lackner. Mass gatherings contributed to early COVID-19 spread: Evidence from US sports. Working Paper 2003, Working Paper, 2020. URL `https://www.econstor.eu/handle/10419/222442`.

Fernando E. Alvarez, David Argente, and Francesco Lippi. A Simple Planning Problem for COVID-19 Lockdown. Working Paper 26981, National Bureau of Economic Research, April 2020. URL `https://www.nber.org/papers/w26981`.

Martin S. Andersen, Ana I. Bento, Anirban Basu, Christopher R. Marsicano, and Kosali Simon. College Openings in the United States Increased Mobility and COVID-19 Incidence. Technical report, July 2021. URL `https://www.medrxiv.org/content/10.1101/2020.09.22.20196048v3`. Type: article.

Hamada S Badr, Hongru Du, Maximilian Marshall, Ensheng Dong, Marietta M Squire, and Lauren 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, November 2020. ISSN 1473-3099. doi: 10.1016/S1473-3099(20)30553-3. URL https://www.sciencedirect.com/science/article/abs/pii/S1473309920305533.

Paolo Bajardi, Chiara Poletto, Jose J. Ramasco, Michele Tizzoni, Vittoria Colizza, and Alessandro Vespignani. Human Mobility Networks, Travel Restrictions, and the Global Spread of 2009 H1N1 Pandemic. *PLOS ONE*, 6(1):e16591, January 2011. ISSN 1932-6203. doi: 10.1371/journal.pone.0016591. URL https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0016591.

Vitaly Belik, Theo Geisel, and Dirk Brockmann. Natural Human Mobility Patterns and Spatial Spread of Infectious Diseases. *Physical Review X*, 1(1):011001, August 2011. doi: 10.1103/PhysRevX.1.011001. URL https://link.aps.org/doi/10.1103/PhysRevX.1.011001.

Paolo Beria and Vardhman Lunkar. Presence and mobility of the population during the first wave of covid-19 outbreak and lockdown in italy. 65:102616, 2021. ISSN 2210-6707. doi: 10.1016/j.scs.2020.102616. URL https://www.sciencedirect.com/science/article/pii/S2210670720308337.

Fred Brauer. Mathematical epidemiology: Past, present, and future. *Infectious Disease Modelling*, 2(2):113–127, May 2017. ISSN 24680427. doi: 10.1016/j.idm.2017.02.001. URL https://linkinghub.elsevier.com/retrieve/pii/S2468042716300367.

Dena Bravata, Jonathan H. Cantor, Neeraj Sood, and Christopher M. Whaley. Back to School: The Effect of School Visits During COVID-19 on COVID-19 Transmission. Working Paper 28645, National Bureau of Economic Research, April 2021. URL https://www.nber.org/papers/w28645.

Philipp Breidenbach and Timo Mitze. Large-scale sport events and COVID-19 infection effects: evidence from the German professional football 'experiment'. *The Econometrics Journal*, (utab021), July 2021. ISSN 1368-4221. doi: 10.1093/ectj/utab021. URL https://doi.org/10.1093/ectj/utab021.

Fiona Burlig, Anant Sudarshan, and Garrison Schlauch. The Impact of Domestic Travel Bans on COVID-19 is Nonlinear in Their Duration. Working Paper 28699, National Bureau of Economic Research, April 2021. URL https://www.nber.org/papers/w28699.

Diego Caccavo. Chinese and Italian COVID-19 outbreaks can be correctly described by a modified SIRD model. Technical report, medRxiv, April 2020. URL https://www.medrxiv.org/content/10.1101/2020.03.19.20039388v3. Type: article.

Brantly Callaway and Tong Li. Policy Evaluation during a Pandemic. May 2021. URL https://arxiv.org/abs/2105.06927v1.

Patrick R. Carlin, Paul Minard, Daniel H. Simon, and Coady Wing. Effects of large gatherings on the COVID-19 epidemic: Evidence from professional and college sports. *Economics & Human Biology*, 43:101033, December 2021. ISSN 1570-677X. doi: 10.1016/j.ehb.2021.101033. URL https://www.sciencedirect.com/science/article/pii/S1570677X21000575.

Riccardo Castagnoli, Martina Votto, Amelia Licari, Ilaria Brambilla, Raffaele Bruno, Stefano Perlini, Francesca Rovida, Fausto Baldanti, and Gian Luigi Marseglia. Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) Infection in Children and Adolescents: A Systematic Review. *JAMA Pediatrics*, 174(9):882–889, September 2020. ISSN 2168-6203. doi: 10.1001/jamapediatrics.2020.1467. URL https://doi.org/10.1001/jamapediatrics.2020.1467.

Jonathan P. Caulkins, Dieter Grass, Gustav Feichtinger, Richard F. Hartl, Peter M. Kort, Alexia Prskawetz, Andrea Seidl, and Stefan Wrzaczek. The optimal lockdown intensity for COVID-19. *Journal of Mathematical Economics*, 93:102489, March 2021. ISSN 0304-4068. doi: 10.1016/j.jmateco.2021.102489. URL https://www.sciencedirect.com/science/article/pii/S0304406821000276.

Serina Chang, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, and Jure Leskovec. Mobility network models of COVID-19 explain inequities and inform reopening. *Nature*, 589(7840):82–87, January 2021. ISSN 1476-4687. doi: 10.1038/s41586-020-2923-3. URL https://www.nature.com/articles/s41586-020-2923-3.

Victor Chernozhukov, Hiroyuki Kasahara, and Paul Schrimpf. The Association of Opening K-12 Schools and Colleges with the Spread of COVID-19 in the United States: County-Level Panel Data Analysis. Technical report, February 2021. URL https://www.medrxiv.org/content/10.1101/2021.02.20.21252131v1. Type: article.

Charles Courtemanche, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz. Strong Social Distancing Measures In The United States Reduced The COVID-19 Growth Rate. *Health Affairs (Project Hope)*, 39(7):1237–1246, July 2020. ISSN 1544-5208. doi: 10.1377/hlthaff.2020.00608.

Charles J. Courtemanche, Anh H. Le, Aaron Yelowitz, and Ron Zimmer. School Reopenings, Mobility, and COVID-19 Spread: Evidence from Texas. Working Paper 28753, National Bureau of Economic Research, May 2021. URL https://www.nber.org/papers/w28753.

Dhaval Dave, Drew McNichols, and Joseph J. Sabia. The contagion externality of a superspreading event: The Sturgis Motorcycle Rally and COVID-19. *Southern Economic Journal*, 87(3):769–807, 2021a. ISSN 2325-8012. doi: 10.1002/soej.12475. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/soej.12475.

Dhaval M. Dave, Andrew I. Friedson, Kyutaro Matsuzawa, Drew McNichols, Connor Redpath, and Joseph J. Sabia. Risk Aversion, Offsetting Community Effects, and COVID-19: Evidence from an Indoor Political Rally. Working Paper 27522, National Bureau of Economic Research, July 2020. URL https://www.nber.org/papers/w27522.

Dhaval M. Dave, Drew McNichols, and Joseph J. Sabia. Political Violence, Risk Aversion, and Non-Localized Disease Spread: Evidence from the U.S. Capitol Riot. Working Paper 28410, National Bureau of Economic Research, February 2021b. URL https://www.nber.org/papers/w28410.

Antonio Desiderio, Giulio Cimini, and Gaetano Salina. Multiplex mobility network and metapopulation epidemic simulations of italy based on open data, 2022. URL http://arxiv.org/abs/2205.03639. type: article.

Yuanyuan Dong, Xi Mo, Yabin Hu, Xin Qi, Fan Jiang, Zhongyi Jiang, and Shilu Tong. Epidemiology of COVID-19 Among Children in China. *Pediatrics*, 145 (6), June 2020. ISSN 0031-4005, 1098-4275. doi: 10.1542/peds.2020-0702. URL https://pediatrics.aappublications.org/content/145/6/e20200702.

Hanming Fang, Long Wang, and Yang Yang. Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-nCoV) in China. Technical Report w26906, National Bureau of Economic Research, March 2020. URL https://www.nber.org/papers/w26906.

Aidan Findlater and Isaac I. Bogoch. Human Mobility and the Global Spread of Infectious Diseases: A Focus on Air Travel. *Trends in Parasitology*, 34(9):772–783, September 2018. ISSN 1471-4922. doi: 10.1016/j.pt.2018.07.004. URL https://www.sciencedirect.com/science/article/pii/S1471492218301429.

Cornelius Fritz and Göran Kauermann. On the interplay of regional mobility, social connectedness and the spread of covid-19 in germany. *Journal of the Royal Statistical Society. Series A,(Statistics in Society)*, 2021.

Germain Gauthier. On the Use of Two-Way Fixed Effects Models for Policy Evaluation During Pandemics. June 2021. URL https://arxiv.org/abs/2106.10949v1.

Hamish Gibbs, Emily Nightingale, Yang Liu, James Cheshire, Leon Danon, Liam Smeeth, Carl A. B. Pearson, Chris Grundy, LSHTM CMMID COVID-19 working Group, Adam J. Kucharski, and Rosalind M. Eggo. Detecting behavioural changes in human movement to inform the spatial scale of interventions against COVID-19. 17(7):e1009162, 2021. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1009162. URL `https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1009162`.

Dan Goldhaber, Scott A. Imberman, Katharine O. Strunk, Bryant Hopkins, Nate Brown, Erica Harbatkin, and Tara Kilbride. To What Extent Does In-Person Schooling Contribute to the Spread of COVID-19? Evidence from Michigan and Washington. Working Paper 28455, National Bureau of Economic Research, February 2021. URL `https://www.nber.org/papers/w28455`.

Andrew Goodman-Bacon and Jan Marcus. Using Difference-in-Differences to Identify Causal Effects of COVID-19 Policies. SSRN Scholarly Paper 3603970, Social Science Research Network, Rochester, NY, May 2020. URL `https://papers.ssrn.com/abstract=3603970`.

Austan Goolsbee and Chad Syverson. Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020. *Journal of Public Economics*, 193: 104311, January 2021. ISSN 0047-2727. doi: 10.1016/j.jpubeco.2020.104311. URL `https://www.sciencedirect.com/science/article/pii/S0047272720301754`.

Sumedha Gupta, Thuy D. Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I. Simon, and Coady Wing. Tracking Public and Private Responses to the COVID-19 Epidemic: Evidence from State and Local Government Actions. Technical Report w27027, National Bureau of Economic Research, April 2020. URL `https://www.nber.org/papers/w27027`.

Jeffrey E. Harris. Geospatial Analysis of the September 2020 Coronavirus Outbreak at the University of Wisconsin – Madison: Did a Cluster of Local Bars Play a Critical Role? Working Paper 28132, National Bureau of Economic Research, November 2020. URL `https://www.nber.org/papers/w28132`.

Hootsuite. Digital in poland: All the statistics you need in 2021, 2022. URL `https://datareportal.com/reports/digital-2021-poland`.

Shankar Iyer, Brian Karrer, Daniel Citron, Farshad Kooti, Paige Maas, Zeyu Wang, Eugenia Giraudy, P. Alex Dow, and Alex Pompe. Large-Scale Measurement of Aggregate Human Colocation Patterns for Epidemiological Modeling. Technical report, December 2020. URL `https://www.medrxiv.org/content/10.1101/2020.12.16.20248272v1`. Type: article.

Benjamin Jeffrey, Caroline E. Walters, Kylie E. C. Ainslie, Oliver Eales, Constanze Ciavarella, Sangeeta Bhatia, Sarah Hayes, Marc Baguelin, Adhiratha Boonyasiri, Nicholas F. Brazeau, Gina Cuomo-Dannenburg, Richard G. FitzJohn, Katy Gaythorpe, William Green, Natsuko Imai, Thomas A. Mellan, Swapnil Mishra, Pierre Nouvellet, H. Juliette T. Unwin, Robert Verity, Michaela Vollmer, Charles Whittaker, Neil M. Ferguson, Christl A. Donnelly, and Steven Riley. Anonymised and aggregated crowd level mobility data from mobile phones suggests that initial compliance with COVID-19 social distancing interventions was high and geographically consistent across the UK. 5:170, 2020. ISSN 2398-502X. doi: 10.12688/wellcomeopenres.15997.1.

Matt J Keeling and Ken T.D Eames. Networks and epidemic models. *Journal of The Royal Society Interface*, 2(4):295–307, September 2005. ISSN 1742-5689, 1742-5662. doi: 10.1098/rsif.2005.0051. URL `https://royalsocietypublishing.org/doi/10.1098/rsif.2005.0051`.

Stephen Kissler, Nishant Kishore, Malavika Prabhu, Dena Goffman, Yaakov Beilin, Ruth Landau, Cynthia Gyamfi-Bannerman, Brian Bateman, Daniel Katz, Jonathan Gal, Angela Bianco, Joanne Stone, Daniel Larremore, Caroline Buckee, and Yonatan Grad. Reductions in commuting mobility predict geographic differences in SARS-CoV-2 prevalence in new york city. 2020. URL `https://dash.harvard.edu/handle/1/42665370`.

Theresa Kuchler, Dominic Russel, and Johannes Stroebel. Jue insight: The geographic spread of covid-19 correlates with the structure of social networks as measured by facebook. *Journal of Urban Economics*, 127:103314, 2022.

Hien Lau, Veria Khosrawipour, Piotr Kocbach, Agata Mikolajczyk, Justyna Schubert, Jacek Bania, and Tanja Khosrawipour. The positive impact of lockdown in Wuhan on containing the COVID-19 outbreak in China. *Journal of Travel Medicine*, 27(taaa037), May 2020. ISSN 1708-8305. doi: 10.1093/jtm/taaa037. URL `https://doi.org/10.1093/jtm/taaa037`.

Wei Lyu and George L. Wehby. Shelter-In-Place Orders Reduced COVID-19 Mortality And Reduced The Rate Of Growth In Hospitalizations. *Health Affairs (Project Hope)*, 39(9):1615–1623, September 2020. ISSN 1544-5208. doi: 10.1377/hlthaff.2020.00719.

Daniel Mangrum and Paul Niekamp. JUE Insight: College Student Travel Contributed to Local COVID-19 Spread. SSRN Scholarly Paper ID 3606811, Social Science Research Network, Rochester, NY, November 2020. URL `https://papers.ssrn.com/abstract=3606811`.

Christopher P. Morley, Kathryn B. Anderson, Jana Shaw, Telisa Stewart, Stephen J. Thomas, and Dongliang Wang. Social Distancing Metrics and Estimates of SARS-CoV-2 Transmission Rates: Associations Between Mobile Telephone Data Tracking and R. *Journal of Public Health Management and Practice*, 26(6):606–612, November 2020. ISSN 1078-4659. doi: 10.1097/PHH.0000000000001240. URL `https://journals.lww.com/10.1097/PHH.0000000000001240`.

Thuy D. Nguyen, Sumedha Gupta, Martin Andersen, Ana Bento, Kosali I. Simon, and Coady Wing. Impacts of State Reopening Policy on Human Mobility. Working Paper 27235, National Bureau of Economic Research, May 2020. URL `https://www.nber.org/papers/w27235`.

OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org , 2017. URL `https://www.openstreetmap.org`.

Katarzyna Orlewska, Waldemar Wierzba, and Andrzej Śliwczynski. Cost-effectiveness analysis of COVID-19 vaccination in Poland. *Archives of Medical Science*, 18(4):1021–1031, December 2021. ISSN 1734-1922, 1896-9151. doi: 10.5114/aoms/144626. URL `https://www.archivesofmedicalscience.com/Cost-effectiveness-analysis-of-COVID-19-vaccination-in-Polanc`

Ján Palguta, René Levínský, and Samuel Škoda. Do Elections Accelerate the COVID-19 Pandemic? Evidence from a Natural Experiment. Working Paper 891, GLO Discussion Paper, 2021. URL `https://www.econstor.eu/handle/10419/235583`.

Valentina Pieroni, Angelo Facchini, and Massimo Riccaboni. COVID-19 vaccination and unemployment risk: lessons from the italian crisis. 11(1):18538, 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-97462-6. URL `https://www.nature.com/articles/s41598-021-97462-6`.

Raquel Pérez-Arnal, David Conesa, Sergio Alvarez-Napagao, Toyotaro Suzumura, Martí Català, Enrique Alvarez-Lacalle, and Dario Garcia-Gasulla. Comparative Analysis of Geolocation Information through Mobile-Devices under Different COVID-19 Mobility Restriction Patterns in Spain. *ISPRS International Journal of Geo-Information*, 10(2):73, February 2021. ISSN 2220-9964. doi: 10.3390/ijgi10020073. URL `https://www.mdpi.com/2220-9964/10/2/73`.

Lisa A. Robinson, Ryan Sullivan, and Jason F. Shogren. Do the Benefits of COVID-19 Policies Exceed the Costs? Exploring Uncertainties in the Age–VSL Relationship. *Risk Analysis*, 41(5):761–770, May 2021. ISSN 0272-4332, 1539-6924. doi: 10.1111/risa.13561. URL `https://onlinelibrary.wiley.com/doi/10.1111/risa.13561`.

Hector Rufrancos, Mirko Moro, and Eve Moore. The impact of university reopenings on covid-19 cases in scotland. *Covid Economics*, 80, 2021. URL `https://cepr.org/sites/default/files/CovidEconomics80.pdfPaper1`.

Harry E. R. Shepherd, Florence S. Atherden, Ho Man Theophilus Chan, Alexandra Loveridge, and Andrew J. Tatem. Domestic and international mobility trends in the united kingdom during the COVID-19 pandemic: an analysis of facebook data. 20(1):46, 2021. ISSN 1476-072X. doi: 10.1186/s12942-021-00299-5. URL `https://doi.org/10.1186/s12942-021-00299-5`.

Evgeniia Shtele, Paolo Beria, and Vardhman Lunkar. Using location based social media data to explain COVID-19 spread in italy. 22(2):132–160, 2022. ISSN 1567-7141. doi: 10.18757/ejtir.2022.22.2.5702. URL `https://journals.open.tudelft.nl/ejtir/article/view/5702`.

Luke Sloan and Jeffrey Morgan. Who tweets with their location? understanding the relationship between demographic characteristics and the use of geoservices and geotagging on twitter. 10(11):e0142209, 2015. ISSN 1932-6203. doi: 10.1371/journal.pone.0142209. URL `https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0142209`.

Alessandro Spelta and Paolo Pagnottoni. Mobility-based real-time economic monitoring amid the COVID-19 pandemic. 11(1):13069, 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-92134-x. URL `https://www.nature.com/articles/s41598-021-92134-x`.

Statcounter. Social media stats poland, 2022. URL `https://gs.statcounter.com/social-media-stats/all/poland/`.

Joseph E. Stiglitz. COVID-19 and Human Freedom | by Joseph E. Stiglitz, September 2021. URL `https://www.project-syndicate.org/commentary/covid19-spike-in-us-reflects-misunderstanding-of-liberty-`

Madison Stoddard, Debra Van Egeren, Kaitlyn E. Johnson, Smriti Rao, Josh Furgeson, Douglas E. White, Ryan P. Nolan, Natasha Hochberg, and Arijit Chakravarty. Individually optimal choices can be collectively disastrous in COVID-19 disease control. *BMC Public Health*, 21(1):832, April 2021. ISSN 1471-2458. doi: 10.1186/s12889-021-10829-2. URL `https://doi.org/10.1186/s12889-021-10829-2`.

Charles A Taylor, Christopher Boulos, and Douglas Almond. Livestock plants and covid-19 transmission. *Proceedings of the National Academy of Sciences*, 117(50):31706–31715, 2020.

Loring J Thomas, Peng Huang, Fan Yin, Xiaoshuang Iris Luo, Zack W Almquist, John R Hipp, and Carter T Butts. Spatial heterogeneity can lead to substantial local variations in covid-19 timing and severity. *Proceedings of the National Academy of Sciences*, 117(39):24180–24187, 2020.

Christopher M. Whaley, Jonathan Cantor, Megan Pera, and Anupam B. Jena. Assessing the Association Between Social Gatherings and COVID-19 Risk Using Birthdays. *JAMA Internal Medicine*, 181(8):1090–1099, August 2021. ISSN 2168-6106. doi: 10.1001/jamainternmed.2021.2915. URL `https://doi.org/10.1001/jamainternmed.2021.2915`.

Jie Xu, Sabiha Hussain, Guanzhu Lu, Kai Zheng, Shi Wei, Wei Bao, and Lanjing Zhang. Associations of Stay-at-Home Order and Face-Masking Recommendation with Trends in Daily New Cases and Deaths of Laboratory-Confirmed COVID-19 in the United States. *Exploratory Research and Hypothesis in Medicine*, pages 1–10, July 2020. ISSN 2472-0712. doi: 10.14218/ERHM.2020.00045. URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7361445/`.

7 Appendix

Figure 12: National traffic

(a) Average Weekly Car Traffic

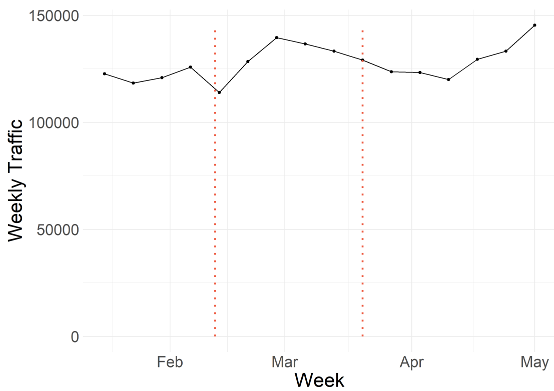

Note: Figure shows average weekly number of passenger cars passing through traffic control points. There are 35 traffic control points equipped with high accuracy cameras distributed throughout the main Polish roads monitoring constantly. For visualization purposes, each point is assigned to the last day of its week (Sunday). Dotted lines represent the start and the end of the hotels opening. Source: Own elaboration based on the data from General Directorate of Roads and Motorways

(b) Monthly Number of Railway Passengers

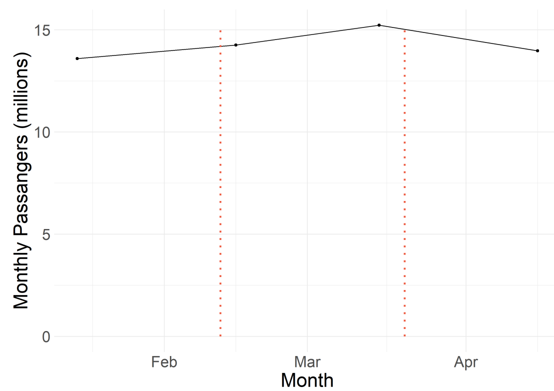

Note: Figure shows the number of passengers transported nationwide by railway in each month. For visualization purposes, each point is assigned to the 15th day of its month. Dotted lines represent the start and the end of the hotels opening. Source: Own elaboration based on data from the Railway Transportation Authority

Table 1: Summary statistics

(a) Population data

Number of unique tiles	34592
Number of tiles with any hotels	4203
Number of tiles of in proximity to ski resorts	3316
Minimum number of users on tile	10
Maximum number of users on tile	8693
Average number of users per tile 01:00-09:00	70
Average number of users per tile 09:00-17:00	81
Average number of users per tile 17:00-01:00	78

Note: Each observation counts the number of FB users on a tile in an 8-hour window. Data is omitted for privacy reasons if there are fewer than 10 users. If a user was on several tiles during the 8-hour period, they are assigned to the tile from which they were logging the most often (modal tile). The period covered is January 6th 2021-March 31st 2021 Source: Facebook Data for Good

(b) Colocation data

Number of unique links	68820
Average number of users with a consistent home per county	768
Average colocation probability	$2.84 * 10^{-6}$
Minimum colocation probability	$9.89 * 10^{-11}$
Maximum colocation probability	0.00156

Note: Each observation corresponds to the probability that two randomly drawn users from two chosen counties meet in a randomly drawn 5-minute interval in a given week. A meeting is defined as being present on the same tile (0.6km × 0.6km) during a 5-minute interval. A user's home county is defined as one where she/he spent at least 6 nights in 10 days intervals around the date considered. User is discarded from computations if there is no consistent night location. Data is omitted for privacy reasons if there are fewer than 10 users. The period considered begins on the third week of January (with the last day 01-09-2021) and ends with the second week of April (with the last day 04-13-2021). Source: Facebook Data for Good

Figure 13: Data summary statistics

(a) Quantiles of FB tile population

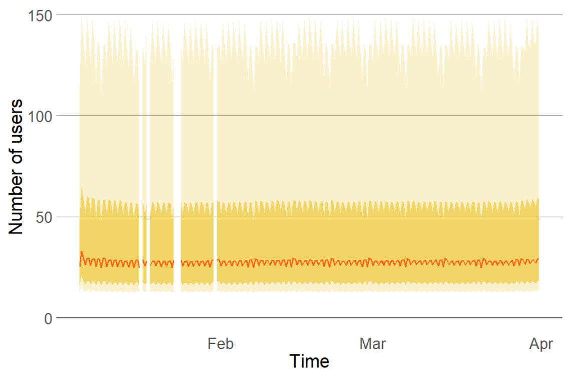

Note: The lighter shaded area corresponds to the 10^{th} and 90^{th} quantiles of the population on tiles in a given 8 hour period. The darker area corresponds to the 25^{th} and 75^{th} quantiles. The red line represents the median. White breaks show missing data. The X-axis shows time which comprises the date and the 8-hour window. Source: Own elaboration based on Facebook data

(b) Quantiles of colocation probabilities

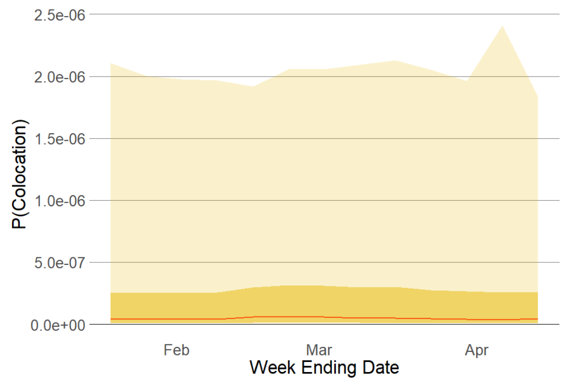

Note: The lighter shaded area corresponds to the 10^{th} and 90^{th} quantiles of colocation probability in a given week. The darker area corresponds to the 25^{th} and 75^{th} quantiles. The red line represents the median. The X-axis shows the last day of the week. Source: Own elaboration based on Facebook data

Facebook data construction and spatial-temporal trends

Facebook geolocation data comes from users who have a Facebook app installed on their phones, and their location history is turned on. Their spatio-temporal records are used to calculate various measures, including Population and Colocation datasets.

Population

Population measures the number of users at a location A in a time-window t. It is computed in the following way.

Assignment of a user to location Map of a country is divided into 3km x 3km tiles. Each 24h is split into three time-periods with breaks 00:00, 08:00, and 16:00 UTC. User is assigned to tile A during time window t if they were pinging from a location on tile A in the time window t. If the user pinged from more than one tile during the time window, they are assigned to their most frequent tile.

Aggregation The population at time window t and at the tile A is the number of users who logged from A in time window t.

Privacy concerns

For privacy reasons, FB does not show data based on fewer than ten users. Given a small size of tiles, there may be many with fewer than 10 users. This is especially true for areas without towns or villages. Mobility in such sparsely populated places may be difficult to estimate. Nonetheless, it is easy to identify when data is missing due to few users present which alleviates the problem.

Spatial-temporal trends Figure 14a shows trends in the daily number of users present in the Population dataset during the study period. The number is relatively stable at about 1 900 000, constituting about 5% of the Polish population. Weekends usually see fewer users than weekdays. There were five days in late January when the number of users was undercounted due to technical difficulties. Nonetheless, these dates are considerably before the policy and do not threaten our strategy. The map on the figure 15a shows the average baseline spatial distribution of the users. The baseline number of users was calculated over 90 days before the data was launched (April 2020). Data relatively well reflects the geographical structure of the Polish population. The tiles with the highest number of users correspond to Poland's large population centers. There are some tiles without enough users to cross privacy threshold in northern Poland which correspond to sparsely populated areas[27].

Colocation

Colocation measures the probability that two randomly chosen users from county r and s were within the same location. It is constructed in the following process.

Assignment of home counties to users. First, a map is divided into administrative units (counties in the case of Poland). Next, each user is assigned a home county based on their nighttime location. Only users who pinged at least three times for each date are counted. The modal location between 8 pm and 6 am is then registered as the user's nighttime location. User is assigned a home in county r if they spent at least 6 out of 10 nights in county r. Let n_r be the number of all users assigned home in county r.

Intersecting users trajectories The goal is to check whether users were at the same place simultaneously. Only users who have location updates sufficiently often are taken into account [28]. A week is divided into 5 minutes intervals, and the map is divided into small 0.6km x 0.6km tiles[29]. Two users met or co-located if their application pinged from the same tile within a 5-minute interval.

Computing colocation measure Let X_{tAr} be the number of users assigned to home county r who pinged on tile A in the 5 minute time interval t. Similarly, let X_{tAs} be the analogous number of users assigned to home county s. Then, the number of collocations produced on tile A at time t between users from counties r and s is the product $X_{tAr}X_{tAs}$. Data is aggregated across all tiles on the map and across all 5 minutes intervals in a week w to compute the number of colocations in the week w: $m_{rs,w} = \sum_{tA} X_{tAr}X_{tAs}$. The colocation probability is then the ratio of all actual meetings and all potential meetings in that week, that is: $Pr(Colocation_{rs,w}) = \frac{m_{rs}}{2016 n_r n_s}$, where 2016 is the number of all five-minute intervals in a week. The procedure is then repeated weekly.

Spatial-temporal trends Figure 14b shows the total number of weekly users (with complete trajectories) available to calculate colocation probabilities. The number varies between 230 000 and 260 000. It is considerably smaller than the Population dataset. This is expected as the requirements to include someone in colocation data are more stringent than for the population data (consistent home, relatively complete trajectories). There is a very slight downward trend during the study period. Map 15b illustrates the average share of county population used to calculate colocation probabilities. These shares are usually between 0.25% and 1% of the total county population. A clear trend arises where a higher share of the population is available in western counties. This follows approximately economic patterns. The divide seems, however, orthogonal to the location of ski resorts.

Usage Predictors

We check whether the main demographic and economic variables correlate with the population share using Facebook geolocation. We pull a set of characteristics at the county level from the Polish Statistical Office for the year 2000. The variables are summarized in table 2. Next, we regress the average share of the

[27] Note that we cannot calculate penetration rate as we do not know true population at a tile level
[28] See Iyer et al. [2020] for technical details
[29] Note that these are smaller tiles compared to population dataset

Figure 14: Temporal trends in Facebook base

(a) Population data

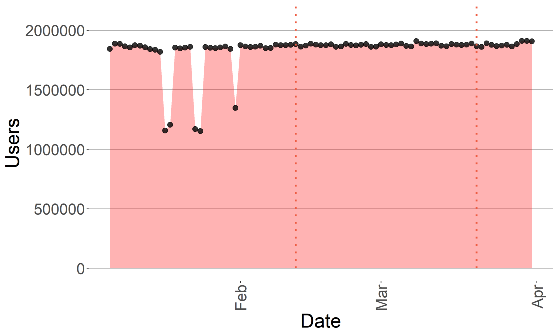

Note: Dots represent the daily number of FB users present in the Population dataset in Poland. Number of users on a given date is an average across three time windows of that day. Vertical dotted lines show the timing of the policy. Note that there were technical issues with data generation on five dates: January 17,18,23,24, and 31.

(b) Colocation data

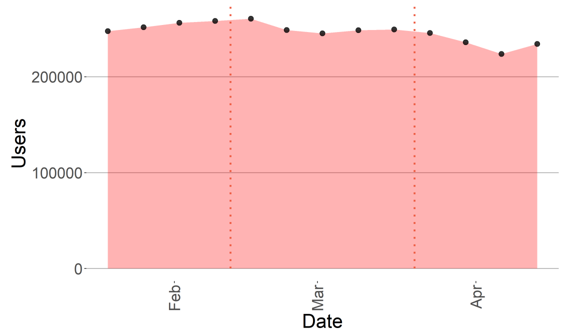

Note: Dots represent the number of FB users in each weak whose data were used to calculate colocation probabilities in Poland. Vertical dotted lines show the timing of the policy.

Figure 15: Spatial trends in Facebook base

(a) Population data

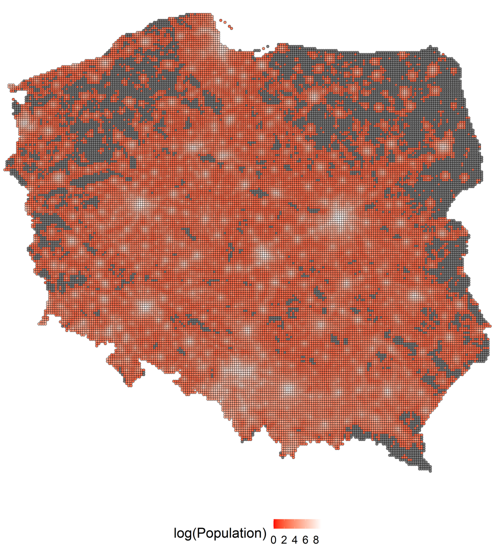

Note: Map shows the average baseline number of FB users on each tile. The baseline is calculated as the mean number of users on a given day-of-the-week and time-window combination across 90 days prior to April 2020 (before pandemic). The map takes average across all day of the weeks and time windows. Grey areas have no data available. The color becomes lighter as log(Population) increases.

(b) Colocation data

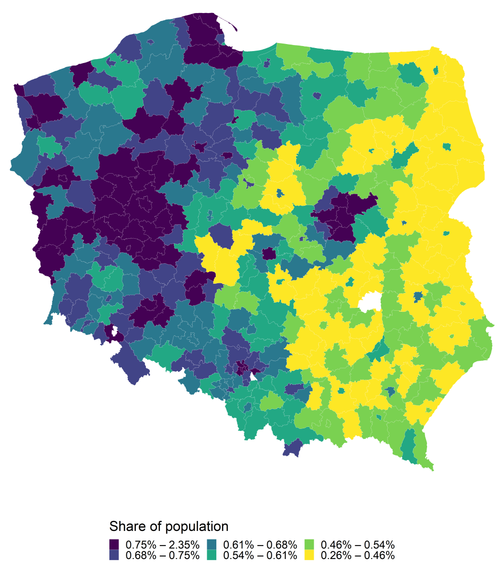

Note: Map shows the average share of county population used to calculate colocation probabilities. The average share in county i is calculated as the mean number of users with home at county i taken across all weeks divided by the total population in this county. The shares are divided into 6 quantiles with dark blue being highest shares and light yellow being lowest shares. White color correspond to no data.

population feeding colocation data (pre-policy) on the county's characteristics. The results are presented in the table 3.

Results suggest that counties with high Facebook geolocation usage are younger, more female and more urban. Moreover, they tend to do better economically as evidenced by the negative coefficient on the share unemployed. On the other hand, they seem to have slightly worse infrastructure in terms of access to healthcare, roads density, and cinemas. Facebook users with geolocation seems to also be more prevalent in counties with more hotel beds. Hence our estimates may put more weight on the movements in these populations. Note, however, that these patterns do not introduce bias in the strategy, as that would require interaction between users' characteristics and the policy timing.

Table 2: Demographic and Infrastructure Variables

Statistic	Mean	St. Dev.	Min	Max
Share in colocation data	0.006	0.002	0.003	0.024
Population	100,920.100	120,344.900	19,689	1,794,166
Population per $1km^2$	362.398	647.498	19	3,690
Share male	0.488	0.009	0.456	0.510
Share age≤14	0.152	0.016	0.110	0.225
Share age≥65	0.182	0.023	0.118	0.284
Share living in urban areas	0.503	0.272	0.000	1.000
Share of university students in the population	0.011	0.031	0.000	0.197
Share unemployed	0.052	0.022	0.016	0.145
Average monthly salary	4,772.427	582.086	3,872.060	8,920.410
Roads per capita	0.009	0.005	0.001	0.033
Doctors per capita	0.002	0.001	0.0001	0.008
Hotel beds per capita	0.025	0.061	0.0002	0.419
Libraries per capita	0.0003	0.0001	0.00003	0.001
Cinema seats per capita	0.005	0.006	0.000	0.033

Impact of opening on colocation

Suppose that the number of tourists who would come from a county s to a county r is proportional to the length of the trails in the county r (LP_r) and the number of people in the county s. Hence, we have $\alpha LP_r n_s$ tourists from s potentially coming to visit r (where α is a proportionality factor). Now, we want to know the number of additional meetings that will occur once the trails are open. In order to have a meeting, individuals need to be in the same space within a five-minute interval. For the moment, assume that every visitor from s to r stays in r for the same amount of time and that space aspect does not matter. That is, suppose that the probability that a tourist meets a local during their stay is δ. So a tourist meets on avarege δn_r locals during their stay. Consequently, the additional number of colocation events is $\alpha\delta LP_r n_s n_r$. Hence, the probability of colocation after the opening is:

$$Pr(Colocation_{rs}|After) = \frac{m_{rs0} + \alpha\delta LP_r n_s n_r}{2016 n_r n_s} = \frac{m_{rs0}}{2016 n_r n_s} + \alpha\delta LP_r$$

where m_{rs0} is the default number of meetings before the opening captured by the fixed effects. Taking logs we have that

$$log(Pr(Colocation_{rs}|After)) = log(m_{rs0} + \alpha\delta LP_r n_s n_s) - log(2016 n_r n_s)$$

Taking the difference between after and before the policy implementation we obtain:

Table 3: Geolocation usage predictors

Dependent Variable:	Share in colocation data
Model:	(1)
Variables	
(Intercept)	0.0292***
	(0.0099)
Population	3.21×10^{-10}
	(8.03×10^{-10})
Population per $1km^2$	-2.63×10^{-7}
	(2.03×10^{-7})
Share male	-0.0513***
	(0.0184)
Share age≤14	0.0242***
	(0.0093)
Share age≥65	-0.0077
	(0.0078)
Share living in urban areas	0.0037***
	(0.0006)
Share of university students in the population	0.0069
	(0.0043)
Share unemployed	-0.0150***
	(0.0038)
Average monthly salary	-8.89×10^{-8}
	(1.49×10^{-7})
Roads per capita	-0.0384
	(0.0248)
Doctors per capita	-0.4296***
	(0.1054)
Hotel beds per capita	0.0028**
	(0.0012)
Libraries per capita	1.369
	(0.8603)
Cinema seats per capita	-0.0337*
	(0.0185)
Fit statistics	
Dependent variable mean	0.00628
R^2	0.38799
Observations	377

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

$$
\begin{aligned}
&log(Pr(Colocation_{rs}|After)) - log(Pr(Colocation_{rs}|Before)) = \\
&log(m_{rs0} + \alpha\delta LP_r n_s n_r) - log(2016 n_r n_s) - \\
&(log(m_{rs0}) - log(2016 n_r n_s)) = \\
&log(\frac{m_{rs0} + \alpha\delta LP_r n_s n_r}{m_{rs0}}) \approx \frac{\alpha\delta n_s n_r}{m_{rs}}
\end{aligned} \tag{6}
$$

Now let us add the hotel beds to the analysis. Assume that the hotel beds attract some additional tourists from county s and that tourists stay longer in ski resorts. In particular, suppose that the number of new visitors is proportional to the number of beds available for them. Hence we have $\tau H_r n_s$ new visitors from s to r (in addition to those who would come just for open trails) where H_r is the number of hotel beds in r and τ is a proportionality constant. Additionally, visitors coming for skiing can now stay longer. Assume again that a share of them proportional to the number of beds stay longer. Hence more meetings can take place. Suppose that the share ζH_r of tourists who stay longer produce κ more meetings than a tourist who does not stay in a hotel. Let us some up all the new terms. First, we have tourists who come skiing but don't stay for the night: $(1 - \zeta H_r)\alpha LP_r n_s$. Second, we have tourists who come because hotels opened: $\tau H_r n_s$. Third, we have tourists who come skiing and stay for the night: $(\zeta H_r)\alpha LP_r n_s$. In total, we obtain the following expression for the colocation probability after the opening of hotels and trails:

$$
\begin{aligned}
Pr(Colocation_{rs}|After) &= \frac{m_{rs0} + (\delta\alpha LP_r n_s n_r)\zeta H_r + \delta\tau H_r n_s n_r + (\kappa\delta\alpha LP_r n_s n_r)(1 - \zeta H_r)}{2016 n_r n_s} \\
&= (\delta\alpha LP_r)\zeta H_r + \delta\tau H_r + (\kappa\delta\alpha LP_r)(1 - \zeta H_r)
\end{aligned} \tag{7}
$$

Taking again the difference of logs before and after the policy we obtain:

$$
\begin{aligned}
&log(Pr(Colocation_{rs}|After)) - log(Pr(Colocation_{rs}|Before)) \approx \\
&\frac{(\delta\alpha LP_r n_s n_r)\zeta H_r + \delta\tau H_r n_s n_r + (\kappa\delta\alpha LP_r n_s n_r)(1 - \zeta H_r)}{m_{rs}}
\end{aligned} \tag{8}
$$

Event studies in the number of Covid-19 Cases Figures 16a and 16b show the coefficients from the health outcomes event study in the full and matched sample, respectively. We estimated the following regression to obtain the coefficients:

$$
y_{kw} = \sum_{\substack{W \in \{\{01/17 : 01/31\}, \\ \{02/14 : 05/02\}\}}} Ski\_resort_k I(w = W)\beta^W + X_{kw}\delta + \zeta_k + \pi_w + \epsilon_{klw} \tag{9}
$$

Where y_{kw} represents the number of cases per 10 000 in a commune k and a week ending at date w. The dummy $Ski\_resorts_k$ takes value 1 if the commune k contains a ski resort, and the indicator $I(w = W)$ is one if the week at hand is equal to W. The interaction between these two terms measures the differential trend in the cases per 10 000 in communes with versus. without ski resorts. X_{kw} contains controls for the share of fully vaccinated two weeks ago and the number of negative tests per 10 000. We allow for time π_w and commune ζ_k fixed effects, and we cluster the errors at the commune level. Figures 16a and 16b plot β_W coefficients from either estimation on the full sample or matched sample. The results are consistent with the unconfoundedness approach and suggest that the opening of hotels sped up the arrival of second-wave to communes with ski resorts.

Figure 16: Event study: Covid-19 cases and hotels opening

(a) Covid cases: full sample

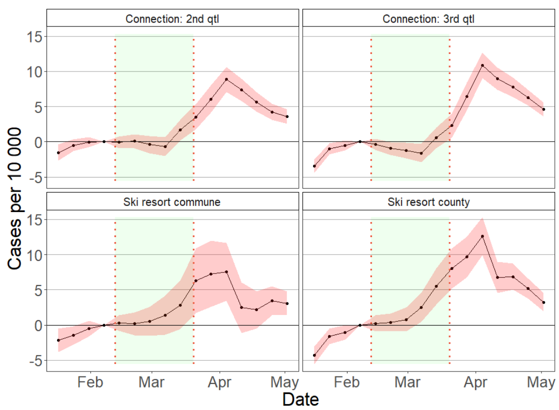

Note: The regression coefficients were calculated on the full sample. The date corresponds to the last day of the week. Source: Own elaboration based on the Ministry of Health Data

(b) Covid cases: matched sample

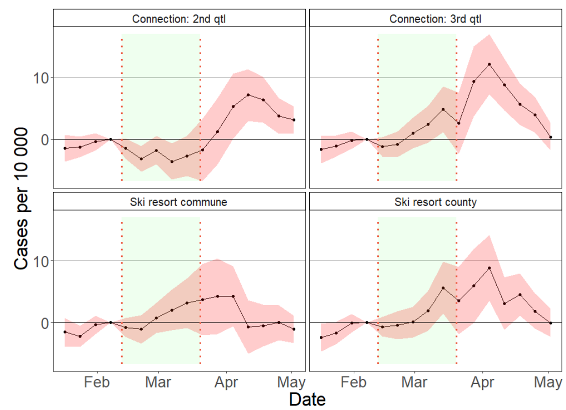

Note: The regression coefficients were calculated on the matched sample. Each treated unit was matched with one untreated units (first quantile of connection). Units were matched by the distance in the propensity scores computed on characteristics in the last period before the policy. The date corresponds to the last day of the week. Source: Own elaboration based on the Ministry of Health Data

Figure 17: Spatial distribution of communes with ski resorts

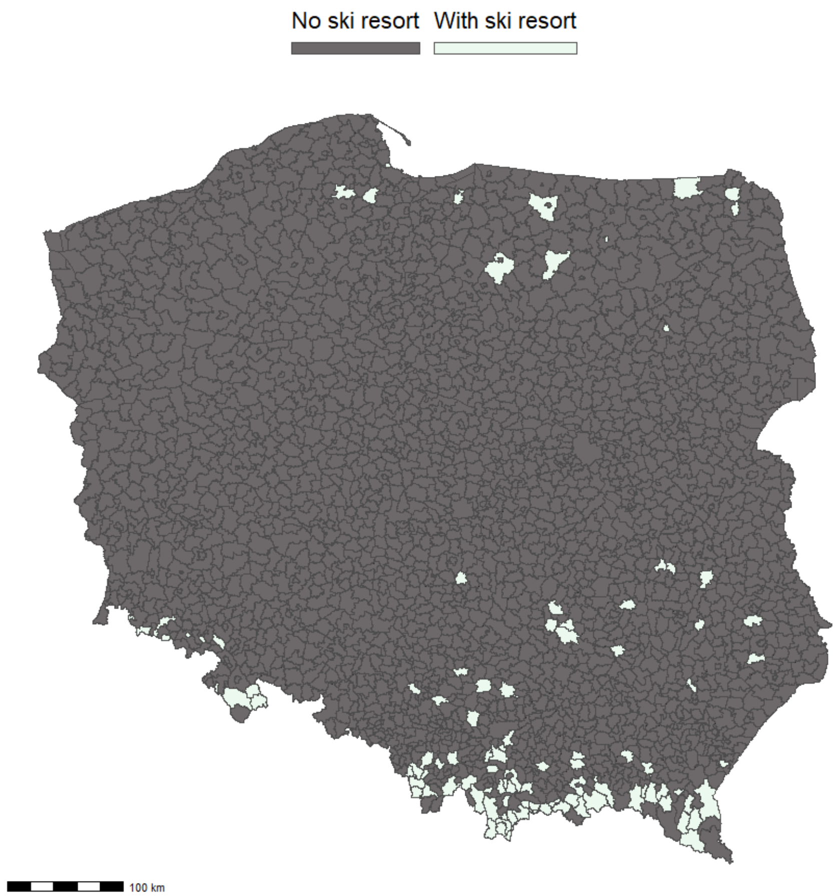

Note: The communes colored in white contain ski resorts. Source: Own elaboration based on data collected from internet

Figure 18: Spatial distribution of the touristic appeal

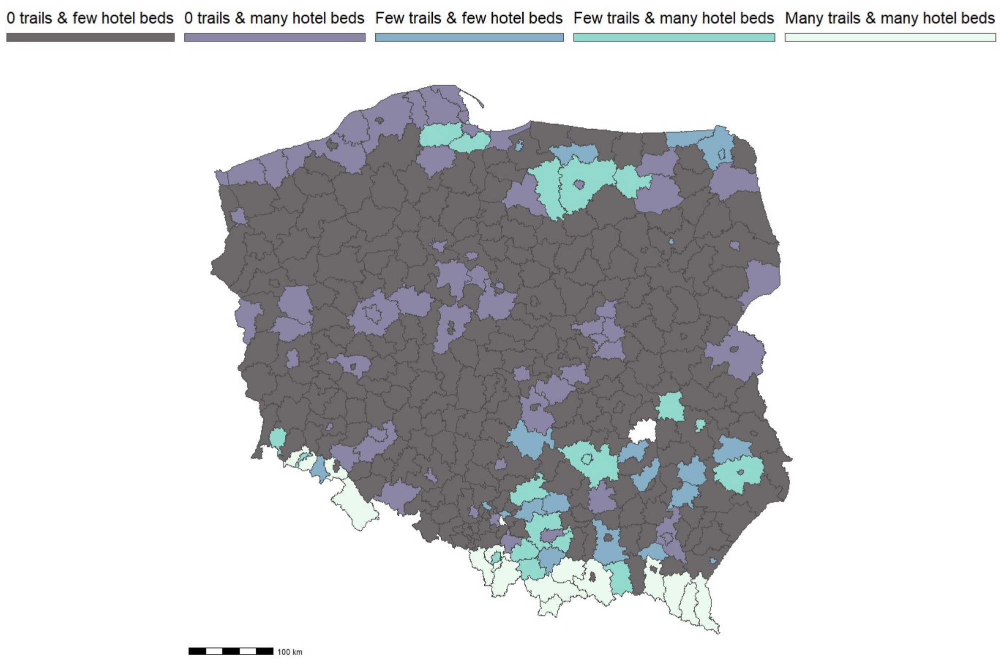

Note: Colors correspond to the touristic appeal. Source: Own elaboration based on data from Polish Statistical Office and own data

Figure 19: Commuting in counties with ski resorets

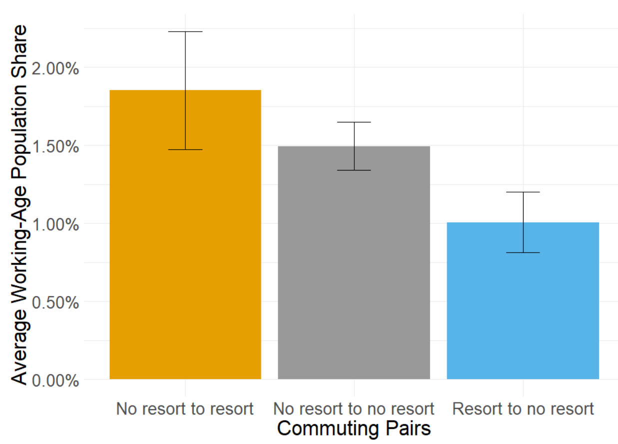

Note: The shares are calculated from data on commuting between communes in 2014 (the most recent available data). A share of working age commuters corresponds to the share of working age population in a commune j commuting to work in a commune i. The averages are taken across types of pairs. No resorts communes are communes without ski resorts and resort communes are communes with ski resorts. The sample is restricted to counties containing ski resorts. Source: Own elaboration based on data from Polish Statistical Office

Figure 20: Event Study: Covid-19 and tourism opening restricted sample

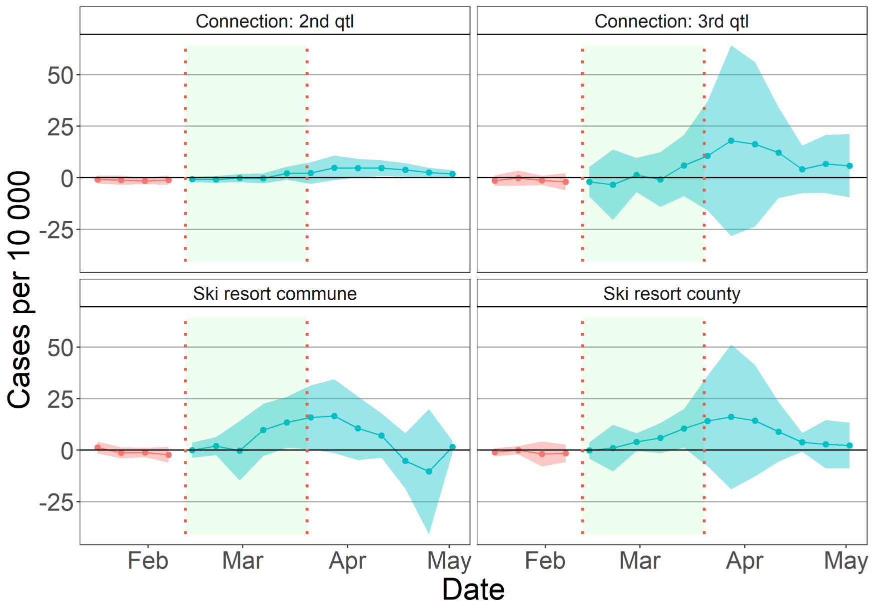

Note: This figure replicates figure 11, but excluding 4 regions which changed restrictions during the study period. As the sample size decreased, some covariates could no longer be used for conditioning due to convergence issues in estimating propensty scores. In partiular, the following covariates were excluded: *population size*, *population density*, and *whether the commune is a county*. The estimates come from the unconfoundedness approach by Callaway and Li [2021]. Red points correspond to the estimates for pre-treatment periods. The reference point in a pre-treatment period t is the previous period $t-1$. Blue points correspond to the estimates for post-treatment periods. The post-treatment periods' propensity scores and outcome regression are based on the last period before the policy $w*-1$. The shaded area represents simultaneous 95% confidence bands with clustering at the commune level. The date corresponds to the last day of the week.

Figure 21: Event Study: Covid-19 and tourism opening; extended pre-period

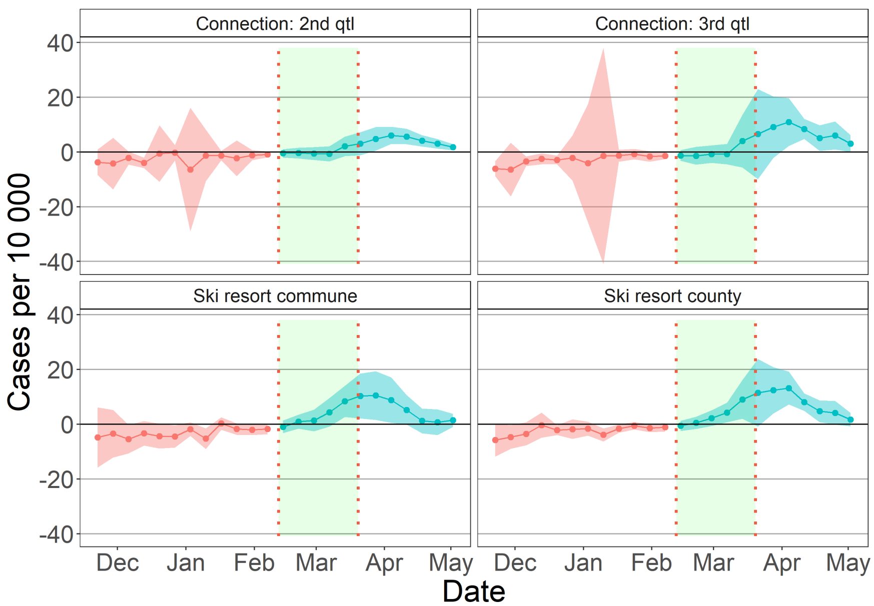

Note: This figure replicates figure 11, but it extends the preperiod until November 2020. Vaccinations and deaths variables can no longer be used for conditioning as the data starts in 2021. The estimates come from the unconfoundedness approach by Callaway and Li [2021]. Red points correspond to the estimates for pre-treatment periods. The reference point in a pre-treatment period t is the previous period $t-1$. Blue points correspond to the estimates for post-treatment periods. The post-treatment periods' propensity scores and outcome regression are based on the last period before the policy $w*-1$. The shaded area represents simultaneous 95% confidence bands with clustering at the commune level. The date corresponds to the last day of the week.

Figure 22: Event Study: Covid-19 and tourism opening; placebo timing

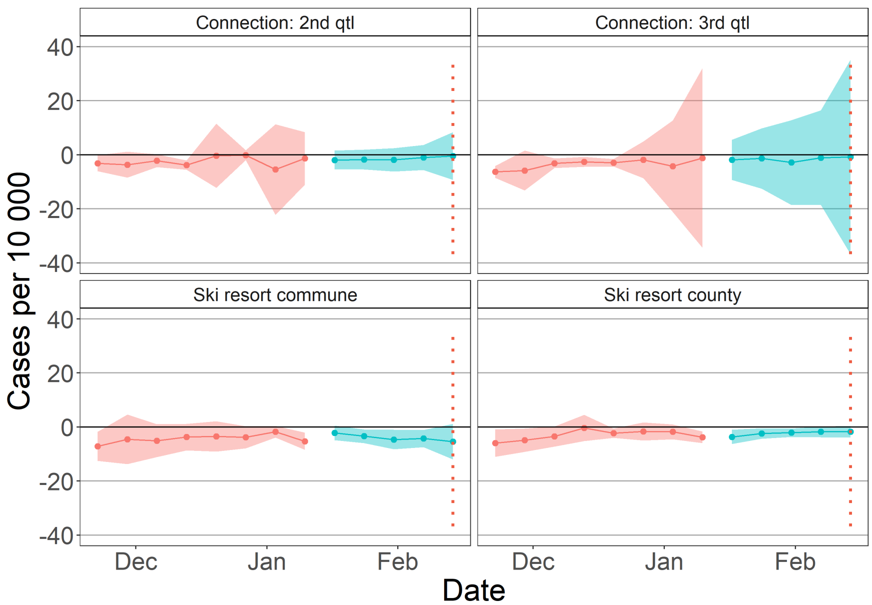

Note: This figure replicates figure 11, but it sets a placebo treatment date on the 14th of January. Vaccinations and deaths variables can no longer be used for conditioning as the data starts in 2021. The estimates come from the unconfoundedness approach by Callaway and Li [2021]. Red points correspond to the estimates for pre-treatment periods. The reference point in a pre-treatment period t is the previous period $t-1$. Blue points correspond to the estimates for post-treatment periods. The post-treatment periods' propensity scores and outcome regression are based on the last period before the policy $w*-1$. The shaded area represents simultaneous 95% confidence bands with clustering at the commune level. The date corresponds to the last day of the week. The dotted line represents the actual start of the actual.