Online Appendix

Regulating Privacy Online: An Economic Evaluation of the GDPR

Samuel G. Goldberg, Garrett A. Johnson, Scott K. Shriver

Table 9: GDPR main effect estimates: panel difference estimator

	(2)
Sample	E-commerce
Dependent variable†	log(Pageviews + 1)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post-GDPR}\}$	-0.118
	(0.030)
$\mathbb{1}\{2018\}$	0.111
	(0.022)
Average marginal effect	-11.15
RSID + Week FE	Y
R^2	0.969
N	22,436

Note: Standard errors clustered at the Dashboard + Week level. †Recorded outcomes.

A Main effect robustness

We first provide pageviews results for the e-commerce sample in Section A.1. To address potential threats to validity, we outline alternative empirical strategies. Section A.2.1 addresses a concern about anticipatory or delayed compliance behavior. Section A.2.2 outlines a synthetic controls approach that relaxes an underlying restriction in our panel difference estimator and allows us to more flexibly construct a control group. Section A.2.3 discusses an alternate control group and two associated model specifications that account for any contemporaneous shocks to the web outcomes of Western users. Section A.3 provides the results for these robustness exercises.

A.1 Pageviews results for the e-commerce sample

Table 9 provides the pageviews outcome estimates for the e-commerce site sample. The panel differences estimation approach mirrors that in our main effect estimates Table 3 and the resulting estimates are comparable to both the full sample pageviews estimate and the e-commerce sample revenue estimate.

A.2 Alternative empirical strategies

A.2.1 Window regressions (PD-WR)

To address concerns about anticipatory or delayed compliance by websites, we remove a two-month window around the GDPR from our data and re-estimate the regression in equation (3). Firms made large investments to comply with the GDPR. Surveys reveal that some firms completed this work before the enforcement deadline while other firm's efforts were ongoing (TrustArc, 2018). Since firms can make quick changes to their website and online marketing, they have an incentive to wait until the last minute to implement their online GDPR compliance changes. Other research confirms that most websites waited until the enforcement deadline before making changes to their sites (e.g. Sørensen and Kosta 2019; Johnson et al. 2022). Note also that Figure 2 displays no change in trend before the GDPR took effect. This is further supported by pre-enforcement placebo checks in Online Appendix A.4. Nevertheless, anticipatory or delayed compliance relative to the deadline may lead us to under- or over-estimate the effect of the GDPR in equation (3). We address this by reestimating equation (3) after dropping four weeks of data both before and after the deadline: we term these the "window regression" results.

A.2.2 Synthetic controls (SC)

Our empirical strategy's within-dashboard design uses the 2017 dashboards as the control group to capture firm-specific seasonality and characteristics. However, the difference-in-difference model assumes that the counterfactual trend is best represented by an equal weighting of these dashboards in 2017. Synthetic controls present a data-driven alternative for selecting the control group. We use the synthetic controls approach (Abadie et al., 2010; Doudchenko and Imbens, 2016) to construct a re-weighted control group that relaxes our within-dashboard restriction in order to predict the counterfactual in the post-GDPR period.

Synthetic controls flexibly construct a control group by taking a weighted average of control-units in order to best predict outcomes for treated units. Intuitively, if we can construct a control group that behaves similarly enough to the treatment group in the pre-period, then this control group should behave similarly to how the treatment group would have behaved after the intervention date, had it not received treatment. We define a *control-unit* to be any 2017 dashboard (logged) data and our *treated-unit* to be the mean of our (logged) 2018 dashboard data, plotted in Figure 7. Because we have one control-unit per dashboard (1,084 in our *full* sample and 353 in our *e-commerce* sample) and only one treated unit, we follow Doudchenko and Imbens (2016) and use an elastic net to construct our synthetic control. Weights are chosen in order to minimize the pre-treatment difference between the treated-unit and potential control units. The intent of synthetic controls is to predict the counterfactual, thus we use cross-validation to incorporate prediction error into our objective function. Then, the counterfactual is constructed by taking the chosen weights and using them to aggregate post-treatment control-unit outcomes. We can then recover the treatment effect by differencing the treated and synthetic control outcomes. Online Appendix A.4.2 details the cross-validation and model fitting procedure.

A.2.3 North American difference-in-differences (DD-NA) and triple panel difference (TPD)

We next consider a contemporaneous control group. Our 2017 control group would not account for any 2018 contemporary confounds like global macroeconomic changes. We thus consider the 2018 web outcomes of North American users from our dashboard sample as an alternate control group. Table 1 reveals that our dashboards have slightly less North American traffic than EU traffic, on average. E-commerce dashboards accrue roughly equal percentages of their revenue from the EU and North America. We explain above that this control group is likely contaminated—owing to within-firm spillovers of GDPR compliance to non-EU users. Thus, any specification with North America as the control group may understate the effect of the GDPR. Other GDPR studies use similar contemporaneous controls (Jia et al., 2021; Aridor et al., 2020; Zhuo et al., 2021), though these authors also acknowledge this contamination issue. We present both our panel estimator with North American controls and a triple panel difference specifications using the North America control. Both specifications address a confounding and contemporaneous shock to online outcomes to both EU and North American users. The triple panel difference specification compares our preferred EU panel difference estimate (equation 3) and an analogous North American panel difference estimate that uses outcomes from 2017 and 2018. The triple panel difference specification identifies the GDPR effect separately from continent-specific seasonality (shared in 2017 and 2018) and a common shock after May 2018 to both EU and North American users.

A.3 Robustness results

In this section, we present the results of our robustness tests. First, in Table 10 we present our panel differences model with alternative fixed effects specifications including (1) a saturated model with dashboard-by-week fixed effects (columns (2) and (4)) and (2) a model that allows for dashboard specific linear time trends. In addition, Table 10 includes the results of our panel differences window regression specification from Section A.2.1. Our main effect results are largely unchanged by the addition of richer fixed effects. Allowing for dashboard-specific, linear-growth trends (columns (3) and (7)) leads to a slight increase (decrease) in our estimated treatment effect for the pageviews (revenue) outcome, though the interpretation and significance of our results remains unchanged. Finally, both our window regression estimates (columns (4) and (8)) are higher than our preferred estimates in Table 3, though they remain within the original confidence intervals.[20] The marginal effects here are -16.0%

[20] To address the possibility of further delayed compliance, we also considered models that instead dropped 5, 6, 7, and 8 weeks of data after the GDPR (as well as 4 weeks before). These alternative specifications had

Table 10: Robustness: alternative fixed effects specifications and window regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	PD - Preferred	PD	PD	PD-WR	PD - Preferred	PD	PD	PD-WR
Sample	All Dashboards				E-commerce			
Dependent variable†	log(Pageviews + 1)				log(Revenue + 1)			
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post-GDPR}\}$	-0.124	-0.124	-0.137	-0.174	-0.142	-0.138	-0.095	-0.176
	(0.026)	(0.026)	(0.022)	(0.025)	(0.030)	(0.03)	(0.028)	(0.031)
$\mathbb{1}\{2018\}$	0.050	0.050		0.056	0.193	0.192		0.215
	(0.014)	(0.014)		(0.015)	(0.028)	(0.027)		(0.028)
Dashboard + Week FE	Y	N	N	Y	Y	N	N	Y
Dashboard x Week FE	N	Y	N	N	N	Y	N	N
Dashboard x Trend	N	N	Y	N	N	N	Y	N
R^2	0.969	0.983	0.738	0.966	0.962	0.983	0.734	0.961
N	69,344	69,344	69,344	49,832	22,436	22,436	22,436	16,111

Note: Standard errors clustered at the Dashboard + Week level. †Columns (1) and (5) are the preferred specifications from Table 3. Columns (2) and (6) present alternative fixed effects specifications. Columns (3) and (7) instead employ dashboard-specific linear time trends. Columns (4) and (8) implement the window regression specification detailed in Section A.2.1

for recorded pageviews and -16.1% for recorded revenue. These higher estimates could arise if several websites delayed their compliance with the GDPR.

Table 11 presents three alternative research designs as described above: difference-in-differences with North America as the control group (denoted by DD-NA), triple differences (TPD), a triple differences specification with rich fixed effects (TPD-FE), and synthetic controls (SC).

Table 11 indicates that our main effect results are robust to several alternate specifications. Both specifications using North America as a control group yield smaller point estimates though each remains negative and statistically significant at the 1% level. Of the two specifications, the more conservative marginal effect estimates are -5.4% for recorded pageviews and -5.6% for recorded revenue. We take these lower point estimates as evidence of spillover effects from the GDPR on North American traffic, as discussed in Section 3. Columns (4) and (8) build off the triple panel differences design in columns (3) and (5) by including week-by-year fixed effects. Finally, our synthetic control results in columns (5) and (10) indicate marginal effects of -8.7% for recorded pageviews though only -1.4% for recorded revenue. Inference using synthetic controls is difficult: we follow placebo procedures outlined

little impact on the window regression results. We thank a reviewer for suggesting this analysis. Results available upon request.

in Abadie et al. (2010). Only 0.2% of placebos achieve a magnitude of prediction error similar to our pageview result. The revenue synthetic controls estimate is less robust, with 22.4% of placebos achieving a similar magnitude of prediction error. In Online Appendix A.4, we discuss the synthetic control procedure and results in detail.

Taken as a whole, these alternate specifications seek to address potential threats to validity in our preferred empirical approach. The results indicate a robust negative impact of the GDPR on recorded site outcomes, though somewhat less so for our revenue results owing to the synthetic controls analysis.

A.4 Placebo tests

We run placebo tests to assess the potential for false positive effect estimates. We implement placebo tests by first choosing a counterfactual treatment week from the pre-GDPR period of our data. Then, equation (19) is estimated using data from before April 25th (pre-GDPR). We exclude one-month before the implementation of the GDPR in order to omit any anticipatory behavior. This procedure is repeated for placebo treatment dates ranging from 14 to 7 weeks prior to May 25th, for a total of 8 placebo tests. These placebo dates are chosen in order to provide adequate pre-trends and post-trends in the data (at least 3 data points before and after the placebo treatment).

$$\log\left(y_{itw}+1\right)=\alpha\mathbb{1}\{2018\}_t+\beta_p\left(\mathbb{1}\{2018\}_t \text{ x } \mathbb{1}\{\text{Post Placebo}\}_w\right)+\theta_i+\eta_w+\epsilon_{itw} \tag{19}$$

The primary coefficient of interest is β_p. Fixed effects are included as in equation (3) and all standard errors are clustered at the dashboard-week level. Significant point estimates are indicative of false positives; a prevalence of false positives may undermine the credibility of our point estimates.

Our placebo results are presented in Figure 6 in grey. Figure 6 also includes our estimated treatment effect in white, to aid in comparison. The placebo results demonstrate the robustness of our identification strategy and point estimates in Table 3. The placebo estimates are smaller in magnitude than our main results and have the opposite sign. All but one estimate—seven weeks prior for the pageviews outcome—are statistically insignificant.Synthetic controls: additional details

A.4.1 Synthetic controls: detailed results

Section A.2.2 discusses the motivation for implementing synthetic controls. Here we discuss the results and methodology behind this approach in more detail.

Table 11: Robustness: alternative research designs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Model[†]	PD - Preferred	DD-NA	TPD	TPD-FE	SC	PD - Preferred	DD-NA	TPD	TPD-FE	SC
Dependent variable[††]	log(Pageviews + 1)					log(Revenue + 1)				
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post-GDPR}\}$	-0.124		-0.067			-0.142		0.028		
	(0.026)		(0.024)			(0.030)		(0.040)		
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post-GDPR}\}\times\mathbb{1}\{\text{EU}\}$			-0.056	-0.055				-0.149	-0.150	
			(0.019)	(0.019)				(0.049)	(0.049)	
$\mathbb{1}\{\text{EU}\} \times \mathbb{1}\{\text{Post-GDPR}\}$		-0.084	-0.036	-0.036			-0.058	0.097	0.097	
		(0.005)	(0.002)	(0.001)			(0.013)	(0.010)	(0.008)	
$\mathbb{1}\{2018\}$	0.050		0.076			0.193		0.050		
	(0.014)		(0.017)			(0.028)		(0.037)		
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{EU}\}$			-0.026	-0.026				0.144	0.144	
			(0.016)	(0.016)				(0.041)	(0.041)	
Average marginal effect	-11.67%	-8.09%	-5.44%	-5.44%	-8.70%	-13.26%	-5.59%	-13.87%	-13.89%	-1.40%
Dashboard + Week FE	Y	N	N	N		Y	N	N	N	
Dashboard + Week + Region FE	N	Y	Y	N		N	Y	Y	N	
Dashboard + Week x Region FE	N	N	N	Y		N	N	N	Y	
R^2	0.969	0.484	0.491	0.491		0.962	0.391	0.397	0.397	
N	69,344	64,931	138,280	138,280		22,436	20,199	42,721	42,721	

Note: Standard errors clustered at the dashboard + week level. [†]Model: 1) DD-NA: North American user outcome control; 2) TPD: triple panel differences; 3) TPD-FE: triple panel differences with week-by-region fixed effects; and 4) SC: synthetic controls. [††]Recorded outcomes.

Figure 6: Main effect panel difference pre-trend placebo tests

(a) Log recorded pageviews (b) Log recorded revenue

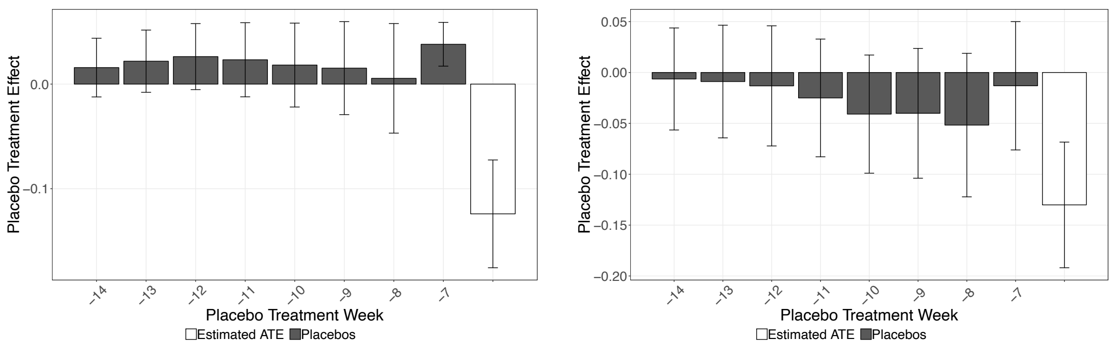

Note: Estimates of placebo treatment effects, β_p from equation (19), with 95% confidence intervals.

Figure 7: Synthetic controls results: fitted trends

(a) Log recorded pageviews (b) Log recorded revenue

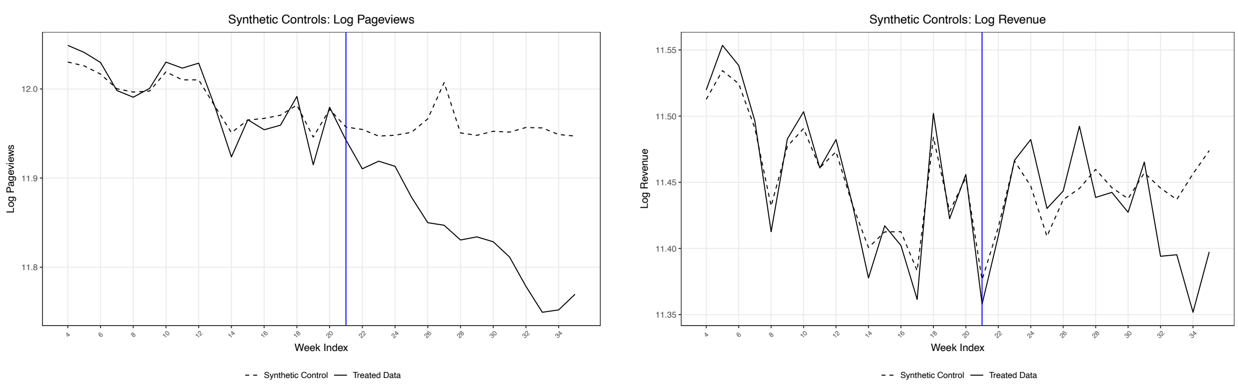

Note: This figure presents the fitted synthetic control groups constructed using the procedure outlined in Online Appendix A.4.

Figure 7 plots average outcomes for our treated group with solid line and our fitted control group in a dotted line. The vertical line marks the implementation of the GDPR. The details of constructing our control group can be found in Section A.4.2. For both pageviews and revenue, Figure 7 illustrates a well fitted control group in the pre-GDPR period. For pageviews, the predicted control group deviates from our 2018 outcomes in the post-GDPR period - congruent with Figure 2. In contrast, our predicted control group continues to match the treated group for revenue well into the post-GDPR period. Our point estimates for each outcome are presented in Table 11. The synthetic controls method estimates a treatment effect of 8.7% for pageviews and 1.4% for revenue.

Constructing standard errors in a synthetic controls setting with an elastic net is not a

Figure 8: Synthetic control results: placebo tests

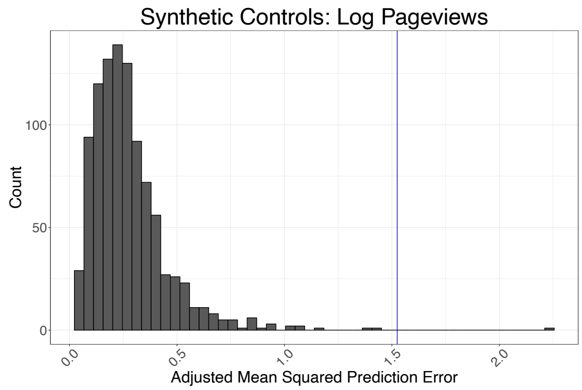

(a) Log recorded pageviews

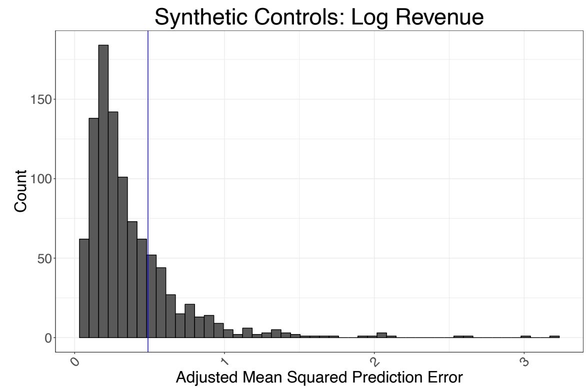

(b) Log recorded revenue

Note: This figure presents the distribution of AMPSE estimates from 1,000 synthetic control placebo studies. Blue lines indicate the true estimated AMPSE.

well understood problem. To provide some sense of the robustness of our synthetic controls results, we appeal to Abadie et al. (2010). Abadie et al. (2010) implement placebo studies by asking "how often would we obtain results of this magnitude if we had chosen a random (counterfactual) treated unit, rather than the factual treated unit?" In our context, we select a placebo treated unit from our control unit, fit a control group using our synthetic controls procedure, and then estimate a placebo treatment effect. We detail this procedure in Section A.4.3. In comparing the placebo and factual synthetic controls, we follow Abadie et al. (2010) in using an adjusted mean squared prediction error statistic (AMSPE).

Figure 8 presents histograms of our 1000 placebo studies for pageviews and revenue. The histograms reflect the distribution of AMSPE estimates across our 1000 placebo trials, for each outcome. The vertical line marks the AMSPE of our true estimated treatment effect, 1.5 for pageviews and 0.49 for revenue. For pageviews, our true AMPSE is at the 99.8th quantile of the placebo AMPSE distribution, which suggests that recovering a treatment effect of the magnitude presented in Table 11 column (4) by chance is highly unlikely. The AMPSE for revenue is at the 78.4th quantile of the placebo AMPSE distribution. Our synthetic controls estimate for revenue is therefore less precise than our pageviews measure, and somewhat more likely to occur by chance.

A.4.2 Cross-validation & estimation

For the following discussion, T_0 will be the period before which the intervention takes place, $Y(0)$ is the counterfactual outcome, and $Y(1)$ is the observed outcome of the treatment. In our setting, we will use 2017 dashboard log outcome data as control units and the average (across dashboards) of 2018 log outcome data as our treatment unit. That is, for each outcome variable, we have a treatment unit and a set of control units, denoted by C.[21]

$$Y_t = \frac{1}{N}\sum_i^N \log\left(y_{it}^{2018}+1\right) \tag{20}$$

$$C = \{C_{it} = \log\left(y_{it}^{2017}+1\right)\forall\, i\} \tag{21}$$

We estimate weights such that the weighted combination of C_{it} best matches Y_t. In our setting, we have 17 pre-treatment time periods and 1084 control units, or $N >> T_0$. We follow Doudchenko and Imbens (2016) in using an elastic net to construct our control group. See Zou and Hastie (2005) for a detailed discussion of elastic nets and their properties. In brief, we fit a model with the following objective function:

$$Q(\mu,\omega|Y_t,C_{it};\alpha,\lambda \text{ for } t<T_0) = ||Y_t-\mu-\omega C_{it}||_2^2+\lambda\cdot\left(\frac{1-\alpha}{2}||\omega||_2^2+\alpha||\omega||_1\right) \tag{22}$$

Where μ is a constant, ω is a vector of length N of weights, and α and λ are penalty parameters chosen by the econometrician.

We choose penalty parameters using a modified version of the cross validation routine proposed in Doudchenko and Imbens (2016). In particular, for a proposed pair of penalty parameters, $\{\alpha',\lambda'\}$, we construct pseudo treated units as follows. First, we partition C into B random partitions of size b. We will refer to a partition as C^b. Each C^b is used to construct a pseudo treated unit, $Y_t^{C^b}$, by taking the average over units $i\in C^b$. We use $\tilde{C} = C\backslash C^b$ as the control units for pseudo treated unit $Y_t^{C^b}$. An elastic net is fitted, using only *pre-intervention* data, to obtain $\{\hat{\mu}^b,\hat{\omega}^b\}$. That is: $\{\hat{\mu}^b,\hat{\omega}^b\} = \text{argmin}_{\mu,\omega}\sum_{t=1}^{T_0}\left(Y_t^{C^b}-\mu-\omega\tilde{C}_{it}\right)^2 + \lambda'\cdot\left(\frac{1-\alpha'}{2}||\omega||_2^2+\alpha'||\omega||_1\right)$. Given the weights estimated above and using the proposed penalty parameters $\{\alpha',\lambda'\}$, we predict the outcome for $Y_t^C(0)$ in $t>T_0$ and construct the mean squared error for each B.

$$Y_t^{C^b}(0) = \hat{\mu}^b+\hat{\omega}^b\tilde{C}_{it} \tag{23}$$

$$CV_B(\alpha',\lambda') = \frac{1}{T-T_0}\sum_{t=T_0}^{T}\left(Y_t^{C^b}(1)-Y_t^{C^b}(0)\right)^2 \tag{24}$$

Model performance is then evaluated using the average of the cross validated mean squared

[21] Note that we are reusing notation in this section. For instance, C refers to the control group here and not the consent potential outcome as elsewhere.

error across our B partitions:

$$CV(\alpha', \lambda') = \frac{1}{B} \sum_b CV_b(\alpha', \lambda') \tag{25}$$

Finally, tuning parameters are chosen such that $\{\alpha, \lambda\} = argmin_{\alpha', \lambda'} CV(\alpha', \lambda')$. Using these tuning parameters, we recover the vector of weights ω needed to construct our synthetic control.

We search over a grid of $\alpha \in [.01, .99]$ in increments of .01 and take advantage of the λ validation built into the *glmnet* package in *R* (Friedman et al., 2010). For each $\{\alpha', \lambda'\}$ we partition the control units into $B = 10$ samples—analogous to 10 cross-fold validation. We repeat the above procedure 100 times for each outcome variables to construct a set of ω vectors, which we average to construct our final weights. We then generate our estimates of the treatment effect by differencing the average levels of the treated unit and the synthetic control in the post-GDPR period.

A.4.3 Synthetic controls: Placebo routine

We can appeal to Abadie et al. (2010) to get a sense of how reasonable our results are. In particular, the following exercise asks: "how large would our prediction error be had treatment not occurred?" To construct this counterfactual, our procedure is as follows:

- Randomly sample $n = 10$ units from C
- Construct our pseudo-treated unit as $C_t^{pseudo} = \frac{1}{n} \sum_{i \in sample} C_{it}$
- Fit an elastic net to C_t^{pseudo} as described in Online Appendix A.4.2
- Calculate the adjusted mean squared prediction error (Abadie et al., 2010):

$$\frac{T_0}{T - T_0} \frac{\sum_{t=T_0}^{T} \left(Y_t^{pseudo}(1) - Y_t^{C(pseudo)}(0) \right)^2}{\sum_{t=0}^{T_0} \left(Y_t^{pseudo}(1) - Y_t^{C(pseudo)}(0) \right)^2} \tag{26}$$

  That is, we calculate the mean squared prediction error and scale it by the mean squared fitting error

- Repeat 1000 times for each outcome

The results of this procedure are presented in Figure 8.

Table 12: Last-touch attribution channel ATE regression estimates

Dependent Variable†	(1) log(Pageviews + 1)	(2) log(Revenue + 1)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Direct Traffic}\}$	-0.089	-0.114
	(0.029)	(0.037)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Search}\}$	-0.077	-0.052
	(0.027)	(0.035)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Display}\}$	-0.307	-0.238
	(0.077)	(0.110)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Email}\}$	-0.109	-0.132
	(0.043)	(0.063)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Social}\}$	-0.244	-0.042
	(0.039)	(0.063)
RSID + Week FE	Y	Y
R^2	0.950	0.937
N	146,095	52,855

Note: Standard errors clustered at the Dashboard + Week level. †Recorded outcomes.

B GDPR effect heterogeneity

B.1 Heterogeneity results

B.1.1 Last-touch attribution regression estimates

We implement the following regression equation to estimate marketing channel-specific GDPR effects:

$$\begin{aligned}\log\left(y_{itwc}+1\right) = &\sum \alpha_c \left(\mathbb{1}\{\text{Post GDPR}\}_w \text{ x } \mathbb{1}\{\text{Channel}\}_c\right) + \\ &\sum \gamma_c \left(\mathbb{1}\{2018\} \text{ x } \mathbb{1}\{\text{Channel}\}_c\right) + \\ &\sum \beta_c \left(\mathbb{1}\{2018\} \text{ x } \mathbb{1}\{\text{Post GDPR}\}_w \text{ x } \mathbb{1}\{\text{Channel}\}_c\right) + \eta_{ic} + \xi_w + \epsilon_{itwc}\end{aligned} \tag{27}$$

where c refers to attribution channel and y_{itwc} denotes recorded outcomes associated with channel c. η_{ic} and ξ_ω are dashboard-channel and week-specific fixed effects, respectively. We interact an indicator for channel c, $\mathbb{1}\{\text{Channel}\}_c$, with each term of equation (3). β_c are the coefficients of interest because they capture the effect of the GDPR by channel c. Under the specified model, the variation identifying each β_c will be differences in outcomes across channels and years in post-May 25 weeks, after accounting for unobservables common to dashboard-channel and week observations, and a common level shift in channel outcomes across years. Table 12 presents the coefficient estimates for the different last-touch attribution channels. Note that these estimates are graphed in Figure 3.

B.1.2 Size heterogeneity results

We examine heterogeneous effects of the GDPR across dashboard size by interacting our main specification (equation 3) with an indicator for large dashboards. This indicator equals one for dashboards with above-median pageviews in the pre-GDPR period. Specifically, we estimate the following regression:

$$
\begin{aligned}
\log\left(y_{itw}+1\right) = \quad & \alpha_1 \mathbb{1}\{2018\}_t + \beta_1 \left(\mathbb{1}\{2018\}_t \text{ x } \mathbb{1}\{\text{Post-GDPR}\}_w\right) + \\
& \alpha_2 \mathbb{1}\{2018\}_t \text{x} \mathbb{1}\{\text{Large Site}\}_i + \gamma_1 \mathbb{1}\{\text{Post-GDPR}\}_w \text{x} \mathbb{1}\{\text{Large Site}\}_i + \\
& \beta_2 \left(\mathbb{1}\{2018\}_t \text{ x } \mathbb{1}\{\text{Post-GDPR}\}_w \times \mathbb{1}\{\text{Large Site}\}_i\right) + \eta_i + \xi_w + \epsilon_{itw}
\end{aligned} \tag{28}
$$

The coefficients of interest are β_1, which estimates the average treatment effect of the GDPR on small sites, and the sum of $\beta_1 + \beta_2$ which is an estimate of the average treatment effect of the GDPR on large sites.Table 13 presents the full regression results examining heterogeneity in site usage.

B.1.3 Regulatory Enforcement Regressions

To construct our index of regulatory enforcement beliefs, we use a European Commissions survey from 2008 of 4,835 data controllers across all countries in the EU European Commission (2008). The survey asks to what extent data controllers agree or disagree with "the data protection law in (OUR COUNTRY) is interpreted and applied more rigorously than in other Member States." Responses are recorded on a four point scale and include a no-response option. We exclude the non-responses and construct a response-weighted average index for each country in the EU that takes values from 0 (all responses are "totally disagree") to 1 (all responses are "totally agree"). We then standardize the index to have mean of zero and variance of one. We additionally utilize data on income (GDP per capita) for each EU country in 2018 from the World Bank, which we expect correlates with advertising and e-commerce revenue.[22]

We explore the role of regulatory strictness empirically, by interacting our panel differ-

[22] https://Data.worldbank.org/indicator/NY.GDP.PCAP.CD?locations=PL-GR-PT-DE-EU

Table 13: Treatment effect heterogeneity by site size

	(1)	(2)	(3)	(4)
Sample	All Dashboard		E-commerce	
Dependent variable†	log(Pageviews + 1)	Pageviews per Visitor	log(Revenue + 1)	Revenue per Visitor
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Large Dashboard}\}$	0.040	-0.050	0.101	-0.329
	(0.034)	(0.130)	(0.053)	(0.187)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\}$	-0.144	0.225	-0.182	0.338
	(0.034)	(0.112)	(0.046)	(0.166)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Large Dashboard}\}$	-0.062	0.176	-0.049	0.399
	(0.024)	(0.124)	(0.046)	(0.188)
$\mathbb{1}\{\text{Post GDPR}\} \times \mathbb{1}\{\text{Large Dashboard}\}$	-0.039	0.132	-0.023	0.426
	(0.020)	(0.112)	(0.050)	(0.179)
$\mathbb{1}\{2018\}$	0.081	-0.379	0.213	-0.167
	(0.018)	(0.113)	(0.043)	(0.177)
Average Marginal Effect				
Small Dashboard	-13.43%	25.77%	-16.69%	40.21%
Large Dashboard	-9.87%	18.79%	-7.85%	0.90%
RSID + Week FE + Size FE	Y	Y	Y	Y
R^2	0.969	0.725	0.960	0.758
N	69,344	69,344	22,812	22,812

Note: Standard errors clustered at the Dashboard + Week level. †Recorded outcomes.

Table 14: Regulatory strictness heterogeneity regressions

	(1)	(2)	(3)	(4)
Dependent Variable[†]	log(Pageviews + 1)		log(Revenue + 1)	
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times$ Strictness	-0.056	-0.040	-0.047	-0.041
	(0.006)	(0.006)	(0.016)	(0.019)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\}$	-0.074	0.041	-0.096	-0.040
	(0.022)	(0.029)	(0.022)	(0.075)
$\mathbb{1}\{2018\} \times$ Strictness	-0.007	-0.008	0.030	0.008
	(0.005)	(0.004)	(0.007)	(0.009)
$\mathbb{1}\{\text{Post GDPR}\} \times$ Strictness	0.032	0.027	0.002	-0.014
	(0.007)	(0.009)	(0.010)	(0.012)
$\mathbb{1}\{2018\}$	0.056	0.049	0.163	-0.025
	(0.008)	(0.024)	(0.015)	(0.059)
$\mathbb{1}\{2018\} \times$ Income		0.019		0.469
		(0.056)		(0.151)
$\mathbb{1}\{\text{Post GDPR}\} \times$ Income		0.101		0.345
		(0.054)		(0.128)
$\mathbb{1}\{2018\} \times \mathbb{1}\{\text{Post GDPR}\} \times$ Income		-0.294		-0.138
		(0.079)		(0.186)
RSID + Week FE	Y	Y	Y	Y
R^2	0.969	0.969	0.962	0.962
N	69,344	69,344	22,812	22,812

Note: Standard errors clustered at the Dashboard + Week level. [†]Recorded outcomes.

ences estimator with our regulatory strictness measure. We estimate the following equation:

$$\begin{aligned}\log\left(y_{itw}+1\right) = {} & \alpha_1 \mathbb{1}\{2018\} + \alpha_2 \mathbb{1}\{\text{Strictness x } \mathbb{1}\{2018\} + \alpha_3 \mathbb{1}\{\text{Income}\} \text{ x } \mathbb{1}\{2018\} + \\ & \gamma_1 \mathbb{1}\{\text{Strictness}\} \text{ x } \mathbb{1}\{\text{Post GDPR}\}_w + \gamma_2 \mathbb{1}\{\text{Income}\}\text{x}\mathbb{1}\{\text{Post GDPR}\}_w + \\ & \beta \mathbb{1}\{\text{Post GDPR}\}_w \text{ x } \mathbb{1}\{2018\} + \\ & \beta_{income} \mathbb{1}\{\text{Income}\} \text{ x } \mathbb{1}\{\text{Post GDPR}\}_w \text{ x } \mathbb{1}\{2018\} + \\ & \beta_{stritcness} \mathbb{1}\{\text{Strictness}\} \text{ x } \mathbb{1}\{\text{Post GDPR}\}_w \text{ x } \mathbb{1}\{2018\} + \eta_i + \xi_w + \epsilon_{itw}\end{aligned} \quad (29)$$

Our interaction coefficient of interest is $\beta_{stritctness}$. Following Jia et al. (2021), we also include an interaction with country-level GDP per capita, because income is a potential confound that correlates with regulatory strictness. Results of this regression are presented below.

C Marginal effects

The models in Sections 3 are non-linear and therefore rely on marginal effects for interpretation. In this section we detail the construction of these marginal effects. Our models are generally of the form:

$$log(y_{itw}+1) = \alpha \mathbb{1}\{2018\}_t + \beta\left(\mathbb{1}\{2018\}_t \text{ x } \mathbb{1}\{\text{Post GDPR}\}_w\right) + \eta_i + \xi_w + \epsilon_{itw} \tag{30}$$

where β captures the GDPR effect. First, we estimate the above regression using the data. Then, for the post-GDPR period, we use our estimates to construct predicted outcomes for both the treated and a counterfactual untreated group:

$$y_{itw}^{untreated} = \exp\left(\alpha + \eta_i + \xi_w + \frac{\sigma^2}{2}\right) - 1 \tag{31}$$

$$y_{itw}^{treated} = \exp\left(\alpha + \beta + \eta_i + \xi_w + \frac{\sigma^2}{2}\right) - 1 \tag{32}$$

Where we have included variances ($Var[\epsilon_{itw}] = \sigma^2$) to account for the expected value of the (log-normal) error terms. These predictions are at the dashboard-week level. We construct the marginal effects using the subsample of dashboard-weeks from 2018 after GDPR enforcement; we denote the set of post-GDPR weeks by W_{post}. We then compute the average marginal effect (AME) as follows:

$$AME = \frac{1}{N}\frac{1}{|W_{post}|}\sum_{i=1}^{N}\sum_{w\in W_{post}}\frac{y_{iw}^{treated} - y_{iw}^{untreated}}{y_{iw}^{untreated}} \tag{33}$$

where $|W_{post}|$ denotes the number of post-GDPR weeks in the sample and N is the number of dashboards in the sample.Testing the privacy frictions mechanism

Websites are reluctant to use obtrusive consent dialogs that add friction to the user's browsing experience. The Irish data regulator notes a "general resistance" among sites to introduce privacy frictions through the consent interfaces (Data Protection Commission, 2020b). For this reason, websites experiment with different consent interfaces to reduce privacy frictions and ensure a high consent rate (Long, 2020). Past research suggests that interruptions can hurt online usage (Lambrecht et al., 2011) and noted the long time required to read websites privacy policies (McDonald et al., 2009). Thus, the specific concern of privacy frictions is that they interrupt the user's browsing and deter users from continuing to browse the site—leading to a decrease in real outcomes. This motivates a simple empirical test: we expect the share of visits where the user *bounces* (browses a single page before leaving) will increase as privacy frictions increase.

We test for a privacy frictions mechanism by examining the effect of GDPR on bounce rates. Bounce rates are defined as the share of site visits with only a single pageview. More generally, bounce rates are a key diagnostic outcome in site analytics. Sites may wish to reduce bounce rates in order to increase ad and e-commerce revenue opportunities. On the other hand, high bounce rates can also indicate that the website effectively communicates information to the user: e.g., the temperature in Paris. In the pre-GDPR period, bounce rates average 42% across all dashboards and 37% in the e-commerce sample.

Bounce rates are useful because they indicate how dashboards seek consent. We expect that bouncing users are unlikely to provide explicit consent, because these users minimally interact with the site. If sites employ opt-out consent, then we expect bounce rates to rise somewhat if the privacy frictions mechanisms holds. If instead the site uses a strict opt-in model, then we expect to see a large drop in bounce rates because we expect few bouncing users will opt in to data recording. Comparing individual dashboard means across the pre- and post-GDPR periods suggests that the large majority of our sites see minimal changes to their bounce rates. In particular, only 3 sites exhibit patterns that are consistent with a strict, opt-in approach in that their bounce rates decrease more than 20 percentage points, but their recorded pageviews fall more than 50%. Thus, the vast majority of websites in our sample appear to employ an opt-out approach for consent, which is consistent with large surveys of website behaviors during that period (Sanchez-Rola et al., 2019; Utz et al., 2019; Johnson et al., 2022). As such, consistent with opt-out consent, we expect bounce rates to rise if privacy frictions are contributing substantially to our point estimates. We use bounce rates to test for the privacy frictions mechanism whereby GDPR consent dialogs dissuade users from further browsing.

We test for a change in bounce rates by reestimating our panel differences model (equation 3) with bounce rates as the dependent variable. Table 4 columns (1) and (2) present the model estimates: -0.275 percentage points (s.e. 0.304) for all sites and -0.354 percentage points (s.e. 0.591) for e-commerce sites. Thus, we find no statistically significant evidence that bounce rates change due to the GDPR and these null effect estimates are precisely estimated. This finding may allay website concerns about the privacy frictions mechanism after the GDPR, though more obtrusive consent dialogs that seek opt-in consent could create greater browsing frictions. We proceed under the assumption that the privacy frictions mechanism does not materially contribute to our estimated GDPR effect.

Table 15: GDPR effect estimates: bounce rate

	(1)	(2)
Sample	All Dashboards	E-commerce
Dependent variable†	Bounce rate (percentage points)	
$\mathbb{1}\{2018\}\times\ \mathbb{1}\{\text{Post GDPR}\}$	-0.275	-0.354
	(0.304)	(0.591)
$\mathbb{1}\{2018\}$	0.311	0.111
	(0.304)	(0.539)
RSID + Week FE	Y	Y
R^2	0.861	0.790
N	69,344	22,436

Note: Standard errors clustered at the Dashboard + Week level. †Recorded outcomes.

D A bounding example

Figure 9 illustrates our bounding approach from Section 4 using a numerical example. Figure 10a presents a simple example of a site that sees its recorded pageviews fall from 3,600 to 3,000 after the GDPR. The fall in recorded pageviews is due to changes in both recorded visits and recorded pageviews per visit. In this example, visits fall from 900 to 600, illustrated by the decrease in the height of the rectangle. As in our data, recorded pageviews per visit, rise post-GDPR from 4 to 5 pageviews. This change is reflected in the increase the rectangles width post-GDPR. Our approach decomposes visits into four types: 1) visits that would remain post-GDPR and provide consent; 2) visits that would remain, but refuse consent; 3) visits that leave post-GDPR, but would consent; and 4) visits that would leave post-GDPR and refuse consent. Fundamentally, we assume that we observe the sum of all four types pre-GDPR, and that the post-GDPR outcomes represent the first type of visits alone. As such, our bounding approach seeks to allocate the difference in recorded outcomes—here, 300 recorded visits and 600 recorded pageviews—to the remaining three user types.

Figure 10b illustrates our bounding approach. The right-hand side of Figure 10b shows the simplest case where the consent effect of the GDPR is maximized under the assumption that the GDPR has no real effect on visits. Consenting users are represented by the lower dark rectangle, while the white square represents non-consenting users. In this case, all the missing 300 visits and 600 pageviews are attributed to non-consenting visits (i.e., type two visits). This implies that non-consenting users are adversely selected, consuming only 2 pageviews per visit. The left-hand side of Figure 10b shows the more complicated case in which the consent effect is minimized and the real effect is maximized. Here, we use the natural lower bound of 1 pageview per visit, which constrains the usage, conditional on visiting, of types two through four. Moreover, Assumption 4 holds that the missing visits, due to the GDPR's real harm, are missing at random with respect to consent. Given these assumptions, our bounds suggest that the pure consent effect represents 200 pageviews (type two visits). The real effect of the GDPR accounts for 400 pageviews: 375 pageviews from type three visits and 25 pageviews from type four visits.

E Bounds estimation

Section 4.1 derives tractable moment conditions, at the level of a dashboard j, which relate empirical moments to bounds on the underlying parameters of interest, $\Theta = \{\underline{\delta},\bar{\delta},\underline{\theta},\bar{\theta}\}$. To begin to estimate these moments, we need empirical equivalents of the quantities $E\left[y_j^{obs}(Z_j)|V_j^{obs}(Z_j)=1\right]$ and $E\left[V_j^{obs}=1|Z_j\right]$ with and without the GDPR, i.e., for $Z \in \{0,1\}$. Also, as we note in Section 4.1, we observe visits rather than the average visit probability $E\left[V_j^{obs}=1|Z_j\right]$. However, note that Corollaries 1 and 2 only require the ratio $\frac{E\left[V^{obs}|Z=1\right]}{E\left[V^{obs}|Z=0\right]}$ and equation (6) suggests that this equals the equivalent ratio for visits:

$$\frac{E\left[V_j^{obs}|Z_j=1\right]}{E\left[V_j^{obs}|Z=0\right]} = \frac{E\left[N_j^{obs}|Z_j=1\right]}{E\left[N_j^{obs}|Z_j=0\right]}$$

A final complication arises because we do not directly observe the usage and visits counterfactual outcomes; rather, we only observe either $Z=0$ or $Z=1$ at a point in time. In particular, simple averages before and after the GDPR are insufficient, as they may conflate seasonal variation with the GDPR's effect. To overcome this, we use the panel differences identification strategy introduced in Section 3. For each dashboard j and each outcome (y^{obs} and V^{obs}), we estimate a GDPR treatment effect β_j via the dashboard level analogue of equation (3). That is,

$$E\left[y_j^{obs}|V_j^{obs}=1,Z_j=0\right] = \frac{1}{T_1}\sum_{t\in T_1} y_{jt}^{2018}$$

$$E\left[y_j^{obs}|V_j^{obs}=1,Z_j=1\right] = \beta_j^y + \frac{1}{T_1}\sum_{t\in T_1} y_{jt}^{2018}$$

$$\frac{E\left[N_j^{obs}|Z_j=1\right]}{E\left[N_j^{obs}|Z_j=0\right]} = \beta_j^N$$

where T_1 denotes the total time observations from the pre-GDPR period of our data.

Then, we use the moment conditions from Corollaries 1 & 2 to estimate Θ. Specifically,

Figure 9: Illustrative example of mechanism identification

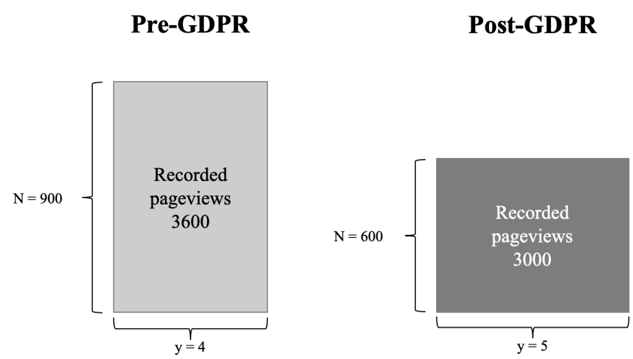

(a) GDPR's effect on a site's recorded outcomes

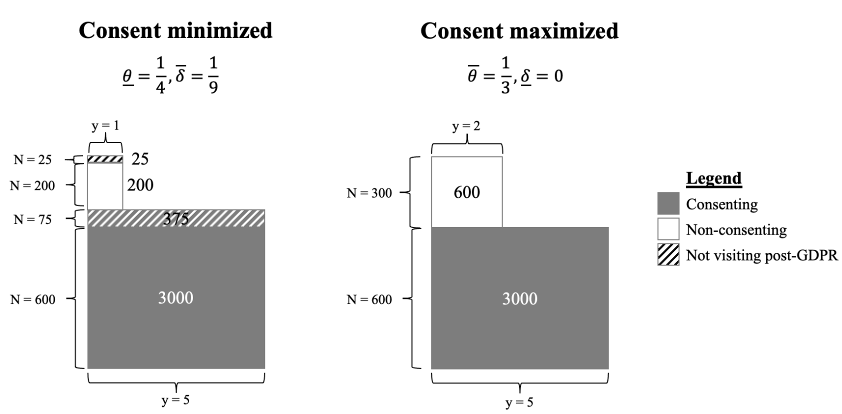

(b) Corresponding bounds on GDPR's consent and real effects

Note: This figure presents a graphical intuition for our bounding analysis in Section 4.

we input the above quantities into equations (12-14) to define our estimation equations:

$$g_{1j}\left(y_j^{2018},\beta_j,\Theta\right) = \underline{\theta}\left(\beta_j^y + \frac{1}{T_1}\sum_{t\in T_1} y_{jt}^{2018} - \underline{y}\right) + \beta_j^y + \varepsilon_{1j}$$
$$g_{2j}\left(y_j^{2018},\beta_j,\Theta\right) = 1 - \bar{\delta} - \frac{\beta_j^N}{(1-\underline{\theta})} + \varepsilon_{2j}$$
$$g_{3j}\left(y_j^{2018},\beta_j,\Theta\right) = 1 - \overline{\theta} - \beta_j^N + \varepsilon_{3j}$$

with the moment condition that $E\left[g\left(y_j^{2018},\beta_j,\Theta\right)\right] = 0$. Note that $\underline{\delta} = 0$ is already determined by equation (15). This yields common Θ estimates across all dashboards, which we estimate using generalized method of moments (GMM). Note that our marketing effect lower bound estimation (Section 4.3) proceeds analogously using equations (16-18).

Conceptually, the ε_j terms relax the requirement that equations (12-14) hold exactly at the dashboard level. We do not directly observe the requisite counterfactual data and our dashboard-level estimates β_j are very imprecise. Instead, we are imposing that these relations hold on average across dashboards. Note that cases where the dashboard level estimates contradict our model—because estimate visits rise or usage falls post-GDPR—are exceptions and are captured by the ε_j terms. The advantage to this approach is that it allows for more heterogeneity in dashboard-level GDPR treatment effects. Another is that we do not impose cross-sectional restrictions on the distributions of dashboard size (i.e., number of visits N_j) and usage metrics (y_j).

F Proofs

Proposition 1

Part A

For part A), we start with the observed usage pre-GDPR $E\left[y^{obs}|V^{obs}=1,Z=0\right]$. Using the potential outcomes notation, we have

$$E\left[y^{obs}|V^{obs}=1,Z=0\right] = E\left[y^{obs}(0)\,|V^{obs}(0)=1,Z=0\right]$$

By the definitions of the recorded outcomes then Assumption 2, we have

$$\begin{aligned} E\left[y^{obs}(0)\,|V^{obs}(0)=1,Z=0\right] &= E\left[y(0)\,|V(0)=1,Z=0\right] \\ &= E\left[y|V(0)=1\right] \end{aligned}$$

Now, we apply the law of total expectation, so that we have

$$\begin{aligned} E\left[y|V(0)=1\right] = & E\left[y|V(0)=1, C(1)=0\right]\Pr\left[C(1)=0|V(0)=1\right] \\ & + E\left[y|V(0)=1, C(1)=1\right]\Pr\left[C(1)=1|V(0)=1\right] \end{aligned}$$

Using Assumption 3, we have

$$\begin{aligned} E\left[y|V(0)=1\right] = & E\left[y|V(0)=1, C(1)=0\right]\Pr\left[C(1)=0\right] \\ & + E\left[y|V(0)=1, C(1)=1\right]\Pr\left[C(1)=1\right] \end{aligned}$$

and by the definition of θ we have

$$\begin{aligned} E\left[y|V(0)=1\right] = & E\left[y|V(0)=1, C(1)=0\right]\cdot\theta \\ & + E\left[y|V(0)=1, C(1)=1\right]\cdot(1-\theta) \end{aligned}$$

Now using Assumption 4, we have

$$\begin{aligned} E\left[y|V(0)=1\right] = & E\left[y|V(0)=1, C(1)=0\right]\cdot\theta \\ & + E\left[y|V(0)=1, V(1)=1, C(1)=1\right]\cdot(1-\theta) \end{aligned}$$

We now show that $E\left[y|V(0)=1, V(1)=1, C(1)=1\right]$ is equivalent to observed usage post-GDPR. By Assumptions 1 then 2, we have

$$\begin{aligned} E\left[y|V(0)=1, V(1)=1, C(1)=1\right] & = E\left[y|V(1)=1, C(1)=1\right] \\ & = E\left[y|V(1)=1, C(1)=1, Z=1\right] \end{aligned}$$

by the definition of recorded outcomes, we have

$$E\left[y|V(1)=1, C(1)=1, Z=1\right] = E\left[y^{obs}|V^{obs}(1)=1, Z=1\right]$$

and simplifying to drop the potential outcomes we have

$$E\left[y^{obs}|V^{obs}(1)=1, Z=1\right] = E\left[y^{obs}|V^{obs}=1, Z=1\right]$$

Putting this all together and expressing in terms of observables, we have

$$\begin{aligned} E\left[y^{obs}|V^{obs}=1,Z=0\right] &= E\left[y\left(0\right)|V\left(0\right)=1,C\left(1\right)=0\right]\cdot\theta \\ &+ E\left[y^{obs}|V^{obs}=1,Z=1\right]\cdot\left(1-\theta\right) \end{aligned}$$

solving for θ obtains the equation in Proposition 1A).

Part B

We start by examining the change in observed visits due to the GDPR:
$E\left[V^{obs}=1|Z=1\right]-E\left[V^{obs}=1|Z=0\right]=E\left[V\left(1\right)\cdot C\left(1\right)=1|Z=1\right]-E\left[V\left(0\right)=1|Z=0\right]$

Where the equality follow from the definition of V^{obs}. Then, given Assumption 3 we have:

$$E\left[V^{obs}=1|Z=1\right]-E\left[V^{obs}=1|Z=0\right]=E\left[V\left(1\right)=1\right]E\left[C\left(1\right)=1\right]-E\left[V\left(0\right)=1\right]$$

Substituting in equations (10) and (9) on the right hand side yields:

$$E\left[V^{obs}=1|Z=1\right]-E\left[V^{obs}=1|Z=0\right]=E\left[V\left(0\right)=1\right]\cdot\left(\left(1-\delta\right)\left(1-\theta\right)-1\right)$$

Finally, we substitute $E\left[V\left(0\right)=1\right]=E\left[V^{obs}=1|Z=0\right]$ and simplify to get:

$$\delta=1-\frac{E\left[V^{obs}=1|Z=1\right]}{E\left[V^{obs}=1|Z=0\right]\left(1-\theta\right)}$$

Proposition 2

Part A

Part A follows from the proof of Corollary 2.

Part B Under Assumption 5, the consent rate estimated in Proposition 2 Part A serves as a lower bound for the Email and Display attribution channels. Then, Part B follows by applying Corollary 1.