Web Appendix
Is Resale Needed in Markets with Refunds? Evidence from College Football Ticket Sales

Drew Vollmer

February 18, 2024

A Application: Screening

A notable feature of refunds is that the primary market seller can offer a menu of contracts to screen consumers based on their future uncertainty. For example, consumers who are less certain about their future demand might be willing to pay a higher price to gain the security of a full refund. Because screening is only possible if different consumers receive different payouts, it can be implemented with refund contracts but not resale. Its effects are therefore important for understanding the relative performance of resale and refunds. However, the analysis so far does not consider screening because there is no heterogeneity in uncertainty or its effects: the consequences and probabilities of idiosyncratic and aggregate shocks H_{ij} and V are the same for all consumers.

In this section, I extend the model to include heterogeneous demand uncertainty and measure the effects of screening with refunds. Heterogeneous demand uncertainty comes from a state of the world realized in the second period. Consumers do not know which state will be realized, and the change in values in each state is heterogeneous.

The empirical motivation is the problem faced by primary ticket sellers during the covid-19 pandemic. When season tickets were sold, it was not clear if there would be a vaccine (the aggregate state ω^{Vax}) or not (the state ω^{NoVax}) by the start of the season. Consumers had varying responses to the state without a vaccine. Some were just as willing to attend without a vaccine, but others would not have paid as much. The variation in penalties without a vaccine is the source of heterogeneity in demand uncertainty. Crucially, if consumers know their preferences for attending without a vaccine when they purchase season tickets, they can sort into different refund contracts in the first period.

The structure of demand uncertainty suggests a simple menu of state-based refund contracts: one offering a full refund when there is no vaccine, one offering a full refund when there is a vaccine, and one without any state-based refund.

The fact that consumer values depend on the aggregate state means that the optimal allocation of tickets is different in the two states. The menu of state-based refunds effects a different allocation in each state, thus affecting total welfare. For that reason, state-based refunds should lead to a more efficient allocation than resale markets. The value of the exercise is to quantify the benefits of refunds. The implications for total welfare also separate the setting from traditional studies of screening on uncertainty, where uncertainty is idiosyncratic and the focus is on an optimal sales mechanism.

Model Changes. There are two possible states of the world in the second period, $\omega \in \{\omega^{\text{Vax}}, \omega^{\text{NoVax}}\}$. The state is unknown in the first period and realized at the start of the second. All consumers expect state ω^{Vax} to occur with probability $\Pr(\omega^{\text{Vax}})$. Because of the new aggregate state ω, I do not allow uncertainty over V in the application.

There is an additional term in consumer utility reflecting the state of the world. Consumer i's utility for tickets of quality q to game j is

$$u_{ijq}(H_{ij}, \omega) = \max\{\ \alpha_j\, (\nu_i - b_i(\omega) + \gamma_q)\, (1 - H_{ij}),\ 0\ \}, \tag{1}$$

where the realized state affects consumer values through the term $b_i(\omega)$. Consumers know their value of $b_i(\omega)$ in the first period, allowing it to influence season ticket decisions.

Assumption 1. *Consumer i knows the value of $b_i(\omega)$ throughout the model. The function satisfies $b_i(\omega^{\text{Vax}}) = 0$, $b_i(\omega^{\text{NoVax}}) \geq 0$, and $b_i(\omega^{\text{NoVax}}) \perp \nu_i$.*

Assumption 1 shows that the term b_i can be interpreted as consumer i's penalty for attending a game when there is no vaccine. Consistent with the penalty interpretation, no consumers have a higher value for tickets when there is no vaccine. The final component of Assumption 1, that penalties and values ν_i are independent, is based on descriptive evidence from the data.

Data and Descriptive Evidence. The extended model requires data to estimate the penalty $b_i(\omega)$ and the probability of each state. In this subsection, I describe the additional data and its main features.

The data come from a survey, conducted in August 2020, asking consumers about their demand for tickets with and without a covid-19 vaccine. Respondents report the maximum they are willing and able to pay for one ticket to a college football game in several scenarios related to covid-19. The scenarios include (i) 2019, as a pre-pandemic benchmark, (ii) a case with a widely available vaccine, and (iii) cases with no vaccine and varying caseloads and social distancing policies. Respondents also report their demographic information and the percent chance of each scenario in September 2021.[1] Eliciting willingness to pay (WTP) and assessments of probabilties by asking directly is used in other surveys (see **?** and **?**).

I distributed the survey to 500 users of Prolific.co, an online distribution platform, in August 2020. Half of respondents were aged 50 or over. The full survey and details can be found in Appendix F.

Although there are several scenarios without a covid-19 vaccine (including ones with and without social distancing), reported willingness to pay is similar in each. The observation that reported values depend mainly on whether there is a vaccine motivated the two-state structure. To combine reported willingnesses to pay from the scenarios without a vaccine to a single no-vaccine state, I average reports across the different caseload scenarios when there is no social distancing, weighting by the average likelihood of each scenario in September 2021 from the survey.

To illustrate consumer heterogeneity, Figure 1 shows reported WTP with a vaccine (the horizontal axis) against the change in WTP from the state with a vaccine to the state without one (the vertical axis).[2] The main conclusion is that there is variation in consumer preferences that can be used for screening. A significant number of consumers have high values in both states of the world (the top right), but many others only have high values when there is a vaccine (the lower diagonal).

A secondary conclusion is that the changes in WTP (similar to $b_i(\omega)$) are not correlated with initial WTP (similar to ν_i). Consumers report proportionately large

[1] The survey also asks about other dates, but only September 2021 is used in the analysis.

[2] The lower triangle is empty because the change in WTP cannot exceed reported WTP.

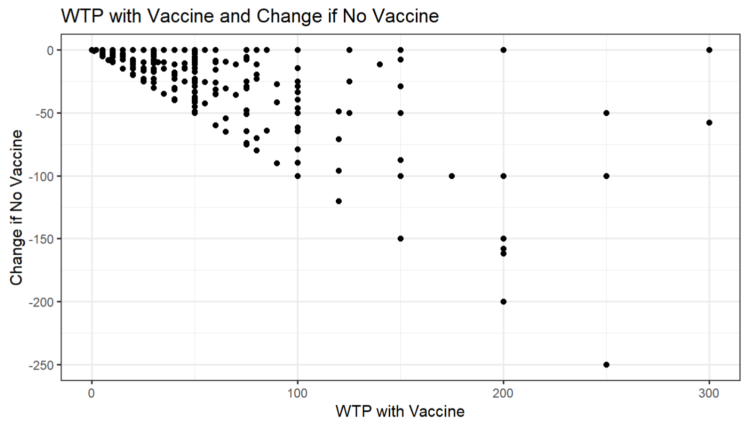

Figure 1: Scatterplot of reported willingness to pay with a vaccine and change in willingness to pay if there is no vaccine. Reported values are for games without reduced seating.

changes in WTP at all levels of initial WTP, and the correlation between the percent change in reported WTP and initial WTP is -.07. The results motivate the independence of $b_i(\omega)$ and ν_i in Assumption 1. Surprisingly, changes in WTP also do not correlate with age. For details, see Appendix F.

Estimation. I use the survey results to estimate two objects: the probability of each state of the world and the distribution of penalties $b_i(\omega^{\text{NoVax}})$.

For the state probabilities, I take the average reported probability from the survey for the states in September 2021, normalized to sum to one. The normalization excludes a case where there is no attendance at sporting events and so no tickets are sold.

For the penalty function, the informative data is the difference in reported willingness to pay between the states with and without a vaccine. For each consumer i, the survey provides observed penalties $\Delta WTP_i = WTP_i(\omega^{\text{Vax}}) - WTP_i(\omega^{\text{NoVax}})$. Assuming that survey responses are for a generic game with game-specific parameter α, and that responses are for the same quality level γ_q, equation (1) implies that the observed differences ΔWTP_i are

$$\Delta WTP_i = \alpha b_i(\omega^{\text{NoVax}}). \tag{2}$$

There are two complications that must be addressed before using equation (2) to estimate the penalty function. First, the parameter α is unknown. I assume that it equals the average of the game-specific parameters α_j, $\bar{\alpha}$. Second, consumer reports of ΔWTP_i are censored: consumers tend not to report negative values of $WTP_i(\omega)$, so the true penalty may be larger than the observed value of ΔWTP_i indicates. Equation (2) must be adjusted to reflect censoring and the use of $\bar{\alpha}$, giving

$$\Delta WTP_i = \min\left\{\bar{\alpha} b_i(\omega^{\text{NoVax}}), WTP_i(\omega^{\text{Vax}})\right\}. \tag{2'}$$

Finally, I assume a parametric form for $b_i(\omega^{\text{NoVax}})$.

Assumption 2. *The function $b_i(\omega^{\text{NoVax}})$ equals zero with probability ρ_1 and is otherwise distributed according to an exponential with parameter ρ_2,*

$$b_i(\omega^{\text{NoVax}}) = \begin{cases} 0 & w.p.\ \rho_1 \\ \tilde{b}_i & \textit{otherwise, where } \tilde{b}_i \sim Exp(\rho_2). \end{cases} \tag{3}$$

I estimate ρ_1 and ρ_2 by maximum likelihood. The likelihood functions are derived from equations (2′) and (3) and are shown in full in Appendix D.6. I calculate standard errors using the bootstrap, repeatedly sampling from the distribution of survey responses.

The estimated state probabilities are shown in Table 1. The parameters defining $b_i(\omega^{\text{NoVax}})$ are shown in Table 2. 29% of consumers have no change in values between states, but the remaining consumers are significantly less willing to attend, reporting a mean penalty of over $52 for the base game.

Table 1: Expected state probabilities in September 2021

State	Probability
Vaccine	0.59
No Vaccine	0.41

Table 2: Estimated preference change parameters. Standard errors calculated using the bootstrap.

Parameter	Value	Std. Err
ρ_1	0.29	0.02
ρ_2	52.27	4.58

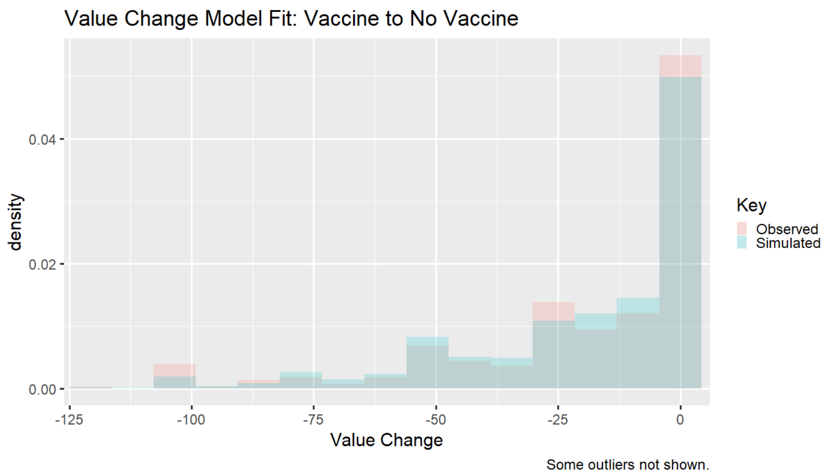

Figure 2: Observed and simulated changes in willingness to pay between states of the world.

Figure 2 shows that observed values of ΔWTP_i are very similar to simulated results. The simulated distribution is not a smooth exponential because, as in the survey, ΔWTP_i is censored at $WTP_i(\omega^{\text{Vax}})$.

One final change is that the screening application uses a different distribution of values ν_i than estimated in Section 6. The reason is that the pandemic may have changed demand for sporting events in September 2021. Therefore, the relevant distribution of values is the distribution of $\nu_{i,Vax}$, not the distribution of $\nu_{i,2019}$ used in the main counterfactuals.

To recover the distribution $\nu_{i,Vax}$, I use survey data on the difference in reported WTP between the 2019 and vaccine states. Using a similar argument to the one used to derive equation (2′), the observed changes in WTP $WTP_{2019} - WTP_{Vax}$ are equal to $\bar{\alpha}(\nu_{i,2019} - \nu_{i,Vax})$. The difference thus looks similar to the penalty function in equation (3). Observed changes are also similar to the earlier penalty, with many responses unchanged between the states but others falling significantly. For that reason, I estimate the difference between $\nu_{i,2019}$ and $\nu_{i,Vax}$ using the same parametric penalty function as in equation (3). Applying the estimated penalty then gives the distribution of $\nu_{i,Vax}$ used below. For details, see Appendix D.

Counterfactual. I consider three aftermarkets: no reallocation, resale, and a menu of state-based refund contracts. The rules of the no reallocation and resale counterfactuals are the same as in the main counterfactuals. The rules of the focal counterfactual are given in Assumption 3.

Assumption 3. *When the primary market seller offers a menu of refund contracts, it prohibits all other ticket transfers. It offers three contracts: a non-refundable package that grants tickets in both states sold at* $\{p^{NR}_{Bq}\}$*, a full refund package granting tickets in the state with a vaccine sold at* $\{p^{FR}_{Bq}(\omega^{Vax})\}$*, and a full refund package granting tickets in the state with no vaccine sold at* $\{p^{FR}_{Bq}(\omega^{NoVax})\}$*. In the second period, the primary seller offers single-game tickets at prices* $\{p_{jq}\}$ *in both states.*

Consumers can only purchase primary market tickets in the second period, and consumers who buy season tickets get value from using their tickets or requesting a refund. As before, the primary market seller maximizes profit by choosing its prices in the counterfactual experiments.

In the counterfactual, I remove uncertainty from idiosyncratic shocks and the common value, $\psi = 0$ and $\sigma^2_V = 0$. The extra sources of uncertainty are not important for measuring the returns to state-dependent contracts and removing them simplifies the results.

Each consumer i is described by a pair $(\nu_i, b_i(\omega^{\text{Vax}}))$ and chooses between the contracts. As in Section 5 of the main text, consumers weigh the surplus from season tickets against the expected surplus from waiting.

Results. The menu of refunds gives the primary seller more control over the final allocation than resale and no reallocation, so it should be more profitable. The contribution of the exercise is to measure the gain in profit and determine the change in welfare.

Table 3 presents the results and demonstrates that the menu of refunds produces significant gains for both consumers and the primary seller. Relative to resale, the menu of refunds boosts total welfare by 6.0%, consumer welfare by 4.5%, and profit by 7.0%.

The partial identification of δ_L has a small effect on the results because a small number of consumers with type L purchase season tickets in the resale counterfactual. I show in Appendix E that the results are virtually unchanged for other values of δ_L in the identified set.

Figure 3 provides evidence on why the menu of refunds performs better. The counterfactual without reallocation performs worst because many consumers with tickets do not want to use them in the state without a vaccine, causing markedly

	No Reall.	Menu of Refunds	Resale
Total Welfare (mn)	9.28	10.01	9.44
	(0.12)	(0.12)	(0.24)
Profit (mn)	6.59	7.23	6.76
	(0.08)	(0.09)	(0.17)
Consumer Welfare (mn)	2.69	2.79	2.67
	(0.04)	(0.04)	(0.10)
Resale Fees (mn)	0.00	0.00	0.01
	(0.00)	(0.00)	(0.00)
Non-Refund. S. Tix (1000)	20.03	12.76	25.87
Vaccine S. Tix (1000)	0.00	5.80	0.00
No Vaccine S. Tix (1000)	0.00	13.07	0.00

Table 3: Counterfactual results for the model with different states of the world. Standard errors calculated using the bootstrap and shown in parentheses.

fewer tickets to be used in that state. Resale and the menu of refunds manage to allocate virtually all tickets in both states, but the menu of refunds avoids the frictions associated with resale.

Figure 4 shows the consumer types in $(\nu_i,\, b_i(\omega^{\text{NoVax}}))$ space choosing each refund contract. The results match the intuition on sorting. Consumers with low penalties (or values so high that the penalty is unimportant) choose the contract without a refund. Consumers with high penalties choose the contract that only grants tickets when there is a vaccine. And consumers with moderate values but relatively low penalties choose the contract that only grants tickets when there is not a vaccine.

B Additional Descriptive Evidence

Figure 5 shows the distribution of normalized prices for the focal university and a random sample of 20 universities. The distributions demonstrate that within-university price variation is significant and widespread. Nearly all universities have a season where prices are 25% above and 25% below the sample mean.

Figure 6 shows seating capacity and how tickets are sold in each zone.

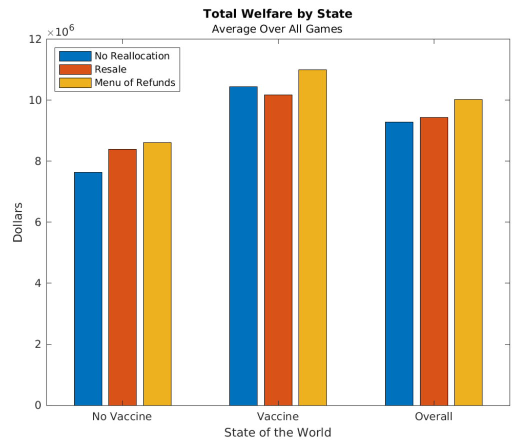

(a) Total welfare.

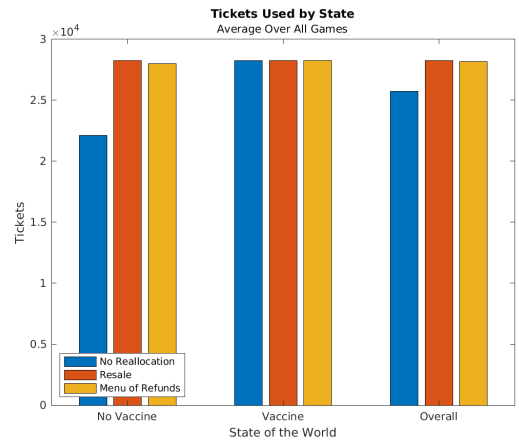

(b) Attendance.

Figure 3: Counterfactual results for total welfare and attendance by state of the world.

C Model Details

Rationing

The random rationing process used in period two is as follows. A consumer starts by requesting his surplus-maximizing choice. She has some chance of receiving the requested choice: probability $\sigma_{jq}(V)$ if she requested a primary market ticket of quality q, or probability one for any other choice. (The resale market clears by Assumption 3 and hence has no stock-outs in equilibrium.) If she does not receive her surplus-maximizing choice, she requests her next-preferred choice. The process ends when a request is accepted.

Consumer i's wait surplus in period one depends on the probability of receiving the surplus-maximizing alternative in her choice set. Let $c^{(m)}(\mathcal{C}_{ij})$ be the m^{th}-largest element of $\mathcal{C}_{ij}$, and let $\sigma_j(V, c)$ be the probability of receiving option c. (The probability $\sigma_j(V, c)$ generalizes the probability $\sigma_{jq}(V)$ of receiving primary market tickets of quality q to game j to any option c.) The expected utility in the second period from waiting with choice set $\mathcal{C}_{ij}$ can be defined recursively as

$$\begin{aligned} \mathit{WaitSurplus}_{ij}(V, \mathcal{C}_{ij}) = \;& \sigma_j(V, c^{(1)}(\mathcal{C}_{ij}))c^{(1)}(\mathcal{C}_{ij}) + \\ (1 - \sigma_j(V, c^{(1)}(\mathcal{C}_{ij})))\, & \mathit{WaitSurplus}_i(V, \mathcal{C}_{ij} \setminus c^{(1)}(\mathcal{C}_{ij})). \end{aligned} \tag{4}$$

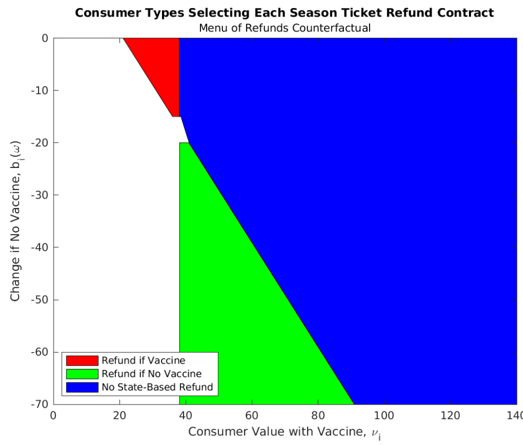

Figure 4: Regions of consumer types selecting each refund contract.

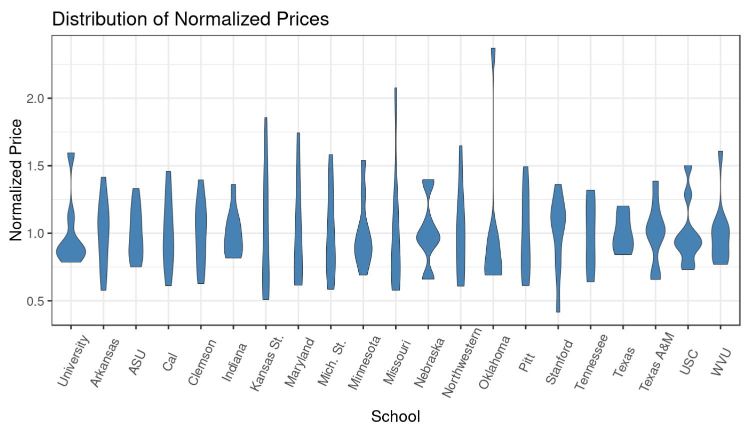

Figure 5: Distribution of average annual resale prices (normalized by school mean) for a random sample of 20 schools in similar conferences.

Overall surplus from waiting is the expected value in the first period,

$$WaitSurplus_i = \sum_j \mathrm{E}_{V,H_{ij},S_{ij}} \left(WaitSurplus_{ij}(V, \mathcal{C}_{ij}(V, H_{ij}, S_{ij})) \right). \tag{5}$$

D Estimation Details

D.1 Complete Pass-Through of Values to Resale Prices

The estimation of several parameters—α_j, γ_q, and σ_V^2—relies on the assumption that changes in values pass through completely to resale prices, as illustrated in equation (11). The assumption holds in equilibrium, as demonstrated by the counterfactual results in Figure 7. Now I elaborate on why the model features complete pass-through.

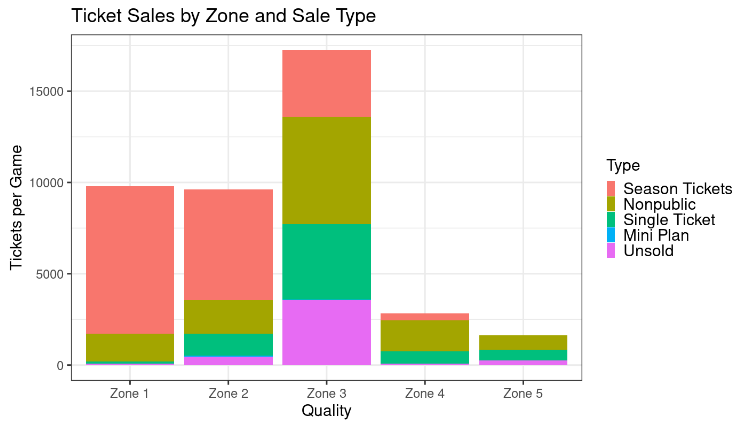

Figure 6: Sale types and volumes by quality group, averaged across games.

The complete pass-through assumption depends on two conditions. First, the supply of tickets available for resale must be roughly fixed. This is plausible in the model because, in equilibrium, only consumers who receive schedule conflicts choose to resell.[3] Second, a shift in demand for all tickets must produce a similar shift in demand for resale tickets. This is plausible if there is limited substitution between the primary and resale markets. Substitution is likely to be limited because resale prices are well below primary market prices in the observed season, but resale frictions imply that there is some substitution.

D.2 Sampling Estimation Moments

This subsection describes the process used to sample estimation moments, which are then used to calculate standard errors of parameters estimated using the method of simulated moments. The overall process is to estimate the variance of each moment using the bootstrap, then take a random sample of moments from the implied distribution. I am unable to consider the covariance of moments when sampling because I only observe one draw of each moment in the data, and the sampling processes do not imply linkages between moments.

Calculating the variance is easiest for each game's resale prices. The data contain records of transaction-level resale prices that can be repeatedly sampled. If there are M_j observed resale transactions for game j, I repeatedly sample M_j draws from the

[3] A caveat is that estimation uses observed (and suboptimal) primary market prices. Thus, in estimation (but not counterfactuals), some consumers who do not receive shocks resell. However, they tend to resell for all games and so the supply of resale tickets remains fixed.

population of transactions and take the variance of the sample average price as the variance for game j.

Calculating the variance is less straightforward for season ticket and primary market quantities because I do not observe consumer-level choices to purchase or not. To handle unobserved choices, I create a data set of simulated choices that matches the data. Specifically, I suppose that there are N total consumers and take N Bernoulli draws with success probability M_j/N, where M_j is the observed number of tickets of the relevant type (season tickets, single-game tickets) purchased. I censor each sample to ensure that no more tickets are sold than are available, then take the variance of the resulting sample means as the moment variance.

One concern with this strategy is that the variance depends on the market size N, which is assumed to be 200,000. If there were no censoring, the variance would follow from N Bernoulli draws with success probability M_j/N,

$$N\frac{M_j}{N}(1-\frac{M_j}{N}) = M_j(1-\frac{M_j}{N}). \tag{6}$$

The only dependence on N is mild because N is large relative to the quantity purchased. Consequently, the last term is close to one and the variance is robust to different values of N.

Moment variances are presented in Table 4.

Table 4: Variance of estimation moments.

Moment	Variance
Season Tickets Sold	19899.16
Avg. Resale Price: Game 1	0.30
Avg. Resale Price: Game 2	0.43
Avg. Resale Price: Game 3	0.31
Avg. Resale Price: Game 4	0.53
Avg. Resale Price: Game 5	0.16
PM Tickets Sold: Game 1	1262.01
PM Tickets Sold: Game 2	3286.64
PM Tickets Sold: Game 3	994.04
PM Tickets Sold: Game 4	2394.55
PM Tickets Sold: Game 5	495.96

D.3 Identifying Season Ticket Parameters

The goal of this subsection is to establish that the season ticket parameters δ_H, δ_L, and ζ are identified. For ease of exposition, assume that all other parameters are fixed at their true values. Start by observing that a candidate value δ_H causes high-type consumers with $\nu_i \geq \nu^*(\delta_H)$ to purchase season tickets. There is an analogous threshold for low types, $\nu^*(\delta_L)$, where $\nu^*(\delta_L) > \nu^*(\delta_H)$.

The most intuitive estimation moment for identification is the number of consumers buying season tickets. For a distribution of ν given by $F_\nu(\nu)$ and market size N, the number of consumers buying season tickets is

$$STQuantitySold = N\left(\zeta\left(1 - F_\nu(\nu^*(\delta_H))\right) + (1-\zeta)\left(1 - F_\nu(\nu^*(\delta_L))\right)\right). \tag{7}$$

Although equation (7) provides a concrete connection between the parameters and the estimation moments, it leaves two free parameters. Additional conditions are needed.

The basis for the additional conditions is that each configuration of parameters defines a unique distribution of values among consumers shopping for tickets in the second period. As long as the estimation moments depend on the full distribution, only the true distribution of second-period values is consistent with observed moments. Denote the distribution by $F_{\nu,2}(\nu)$. Its CDF is

$$F_{\nu,2}(\nu) = \begin{cases} 1 & \text{if } \nu > \nu^*(\delta_L) \\ \left((1-\zeta)F(\nu) + \zeta F(\nu^*(\delta_H))\right)/STFrac & \text{if } \nu^*(\delta_H) < \nu \leq \nu^*(\delta_L) \\ F(\nu)/STFrac & \text{if } \nu \leq \nu^*(\delta_H), \end{cases} \tag{8}$$

where the fraction of consumers who buy season tickets, $STFrac = (1-\zeta)F(\nu^*(\delta_L)) + \zeta F(\nu^*(\delta_H))$, normalizes the CDF.

If any parameter deviates from its true value, then the empirical analogue $\hat{F}_{\nu,2}(\nu)$ will be incorrect on at least one interval. I show that two estimation moments, resale prices and the quantity of tickets sold in the primary market, depend on the distribution $F_{\nu,2}(\nu)$.

The number of primary market tickets sold in quality $\bar{q}$ for game j is

$$
\begin{aligned}
PMTicketsSold_{j\bar{q}} &= (N - STQuantitySold) \Pr(Consumer\ i\ buys\ PM\ ticket\ for\ j,\ \bar{q}| \\
&\quad Did\ not\ buy\ at\ t=1) \\
&= (N - STQuantitySold) \prod_{\{q: u_{ijq}(V) - p_{jq} > u_{ijq}(V) - p_{j\bar{q}}\}} (1 - \sigma_{jq}(V))\sigma_{j\bar{q}}(V) \\
&\quad \Pr(\mathbb{I}\left[s_{ij} > p_{j\bar{q}} - p^r_{j\bar{q}}(V)\right] \mathbb{I}[u_{ij\bar{q}}(V) \geq p_{j\bar{q}}] | Did\ not\ buy\ at\ t=1) \\
&= (N - STQuantitySold) \prod_{\{q: u_{ijq}(V) - p_{jq} > u_{ijq}(V) - p_{j\bar{q}}\}} (1 - \sigma_{jq}(V))\sigma_{j\bar{q}}(V) \\
&\quad \Pr(s_{ij} > p_{j\bar{q}} - p^r_{j\bar{q}}(V)) \left(1 - F_{\nu,2}(\frac{1}{\alpha_j} p_{j\bar{q}} - V - \gamma_{\bar{q}})\right).
\end{aligned}
$$

With sufficient variation in primary market prices p_{jq} relative to the game-specific parameters α_j, the total number of primary market tickets sold—the sum of the $PMTicketsSold_{jq}$ over q—depends on the full distribution of values in the second period.

A similar argument holds for resale prices. Let the function *ResaleSupply* capture the supply side of the resale market (which does not depend on the season ticket parameters). The resale price function $p^r_{j\bar{q}}(V)$ solves the equilibrium condition

$$
\begin{aligned}
ResaleSupply_{j\bar{q}}(p^r_{jq}(V), V) &= ResaleDemand_{j\bar{q}}(p^r_{jq}(V), V) \\
&= (N - STQuantitySold) \cdot \Pr(Consumer\ i\ buys\ resale \\
&\quad ticket\ for\ j,\ \bar{q}|Did\ not\ buy\ at\ t=1) \\
&= (N - STQuantitySold) \int \prod_{\{q: p_{jq} \leq p^r_{jq}(V) + s\}} (1 - \sigma_{jq}(V)) \\
&\quad \Pr(u_{ij\bar{q}} \geq p^r_{jq}(V) + s | s, Did\ not\ buy\ at\ t=1)\, dF_s(s) \\
&= (N - STQuantitySold) \int \prod_{\{q: p_{jq} \leq p^r_{jq}(V) + s\}} (1 - \sigma_{jq}(V)) \\
&\quad \Pr(\alpha_j \nu - s \geq p^r_{jq}(V) - \alpha_j(V + \gamma_q)| \\
&\quad s, Did\ not\ buy\ at\ t=1)\, dF_s(s).
\end{aligned}
$$

The expression depends on $F_{\nu,2}(\nu)$ through the probability term in the integral.

Variation in the frictions s make it so that equilibrium resale prices depend on a wide range of the support of ν in the second period.

D.4 Counterfactual Standard Errors

Standard errors for counterfactual outcomes are calculated using the bootstrap. I obtain the distribution of counterfactual outcomes by running the counterfactual experiments for samples from the distribution of model parameters. The distribution of optimal parameters is taken from the optimal parameters obtained when calculating parameter standard errors in Section 6.

When sampling from the distribution of parameters, I am forced to assume that some parameters are uncorrelated with each other. The reason is that the estimation procedures for some parameters are completely separate, making it impossible to determine their covariance. For example, the parameter λ_s is estimated using the method of simulated moments, but the set of α_j are estimated based on equation (11) and there is no link between the two processes.

The sampling procedure accounts for covariance between parameters whenever possible. For example, the parameters α_j and γ_q are estimated jointly and so their draws are correlated. The same is true of the parameters λ_s, λ_ν, δ_H, δ_L, and ζ, all estimated using the method of simulated moments.

D.5 Profit-Maximizing Prices

Whether profit-maximizing prices in the model are close to observed prices may be useful for evaluating the model fit. Note, however, that the model does not attempt to rationalize observed prices. The estimation procedure only explains demand at observed prices; profit-maximizing prices are not considered or calculated in estimation. Instead, optimal prices are taken from the counterfactual experiment where the university allows resale. No model features ensure that the model's optimal prices are close to those observed in the data.

Table 5 shows model-implied profit-maximizing primary market prices for single-game tickets and season tickets. To easily compare the model-implied and observed values, Table 6 shows model-implied prices minus observed prices.

A broad reading of Table 6 is that observed prices are a bit too low on average, and that they do not vary enough across games. Games 2 and 4 are far too cheap

Table 5: Model-implied profit-maximizing primary market prices for each game.

Game	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
1	\$64.87	\$52.82	\$47.29	\$42.22	\$42.92
2	\$108.04	\$87.97	\$78.76	\$70.31	\$71.48
3	\$65.38	\$53.24	\$47.67	\$42.55	\$43.26
4	\$103.91	\$84.61	\$75.75	\$67.62	\$68.75
5	\$36.59	\$29.79	\$26.67	\$23.81	\$24.21
Season Tickets	\$315.80	\$246.67	\$214.95	\$185.82	\$189.84

Table 6: Model-implied profit-maximizing primary market prices minus observed primary market prices.

Game	Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
1	−5.13	−7.18	−2.71	2.22	12.92
2	38.04	27.97	23.76	25.31	41.48
3	−4.62	−6.76	−2.33	2.55	13.26
4	33.91	24.61	20.75	22.62	38.75
5	−23.41	−25.21	−13.33	−11.19	−5.79
Season Tickets	0.80	−23.33	−1.05	6.82	64.84

compared to the model's profit-maximizing prices, and Game 5 is far too expensive. The university also tends to set prices too low in zone 5.

D.6 Vaccine Demand

Estimating ρ_1 and ρ_2. For the likelihood function used to estimate ρ_1 and ρ_2, let $G(\Delta WTP; \rho_2)$ be an exponential CDF with density $g(\cdot)$. The likelihood function can be written as

$$\mathcal{L}(\rho_1, \rho_2) = \mathcal{L}_1(\rho_1)\mathcal{L}_2(\rho_2), \tag{9}$$

$$\mathcal{L}_1(\rho_1) = \prod_{i=1}^{N} \left(\rho_1^{\mathbb{I}[\Delta WTP_i=0]}(1-\rho_1)^{\mathbb{I}[\Delta WTP_i>0]}\right)^{\mathbb{I}[WTP_i(\omega^{\text{Vax}})>0]}, \tag{10}$$

$$\begin{aligned}\mathcal{L}_2(\rho_2) = \prod_{i=1}^{N} \Big(& g(\Delta WTP_i; \rho_2)^{\mathbb{I}[\Delta WTP_i < WTP(\omega^{\text{Vax}})]} \cdot \\ & (1 - G(\Delta WTP_i; \rho_2))^{\mathbb{I}[\Delta WTP_i = WTP(\omega^{\text{Vax}})]}\Big)^{\mathbb{I}[\Delta WTP_i>0]}.\end{aligned} \tag{11}$$

The likelihood function can be separated into two pieces, one evaluating the likelihood of observing zero change in willingness to pay ($\mathcal{L}_1$) and the other evaluating the likelihood of a positive change ($\mathcal{L}_2$). The function $\mathcal{L}_2$ accounts for censoring: when the observed change ΔWTP_i is the same as initial willingness to pay $WTP(\omega^{\text{Vax}})$, the exact value of the penalty b_i is not observed. Instead, we only know that it is higher than initial willingness to pay, and in such cases the likelihood must be evaluated using the CDF. The case where $WTP_i(\omega^{\text{Vax}})$ equals zero does not affect the likelihood because ΔWTP_i then equals zero with probability one.

Standard errors are calculated using the bootstrap, where samples are drawn from the distribution of survey responses.

Adjusting the Distribution of Values. Recall from Section 6 that the estimated distribution of values from structural estimation, parameterized by λ_ν, reflects demand before covid-19. The survey results suggest that demand with a vaccine is different, as illustrated in Figure 9. Accordingly, I use the survey data to adjust the distribution of ν_i to reflect post-pandemic demand.

Assuming that reported WTP is for a representative game and quality and that there are no relevant penalties, equation (3) implies that the observed differences in WTP from 2019 to the vaccine state are

$$WTP_{i,2019} - WTP_{i,Vax} = \nu_{i,2019} - \nu_{i,Vax}. \tag{12}$$

For that reason, it is appropriate to adjust the distribution of ν_i in the screening application. The data suggest that there is no change in WTP for many consumers, but there is a long tail of consumers who report lower values. The pattern resembles the changes in WTP for the transition in values from the state with a vaccine to the state without one. Therefore, I model the difference as a penalty b'_i using the parametric form in equation (3). The new distribution of values follows

$$\nu_{i,Vax} = \nu_{i,2019} - b'_i, \tag{13}$$

where b'_i is parameterized by ρ_1^{2019} and ρ_2^{2019}. The estimation procedure is the same as for the original penalty function $b(\omega^{\text{Vax}})$; estimation is by maximum likelihood and standard errors are obtained using the bootstrap.

The results are shown in Table 7. Compared to the transition between states with and without a vaccine, many more consumers experience no change in values, 60% as

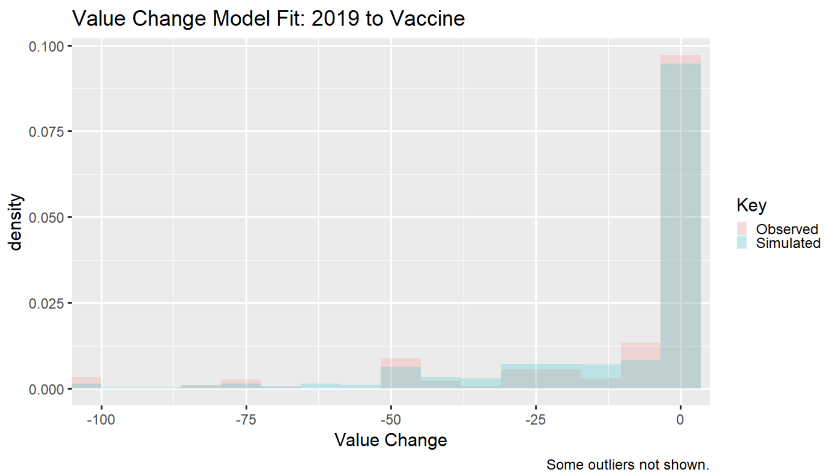

Figure 7: Observed and simulated changes in willingness to pay from 2019 to the state with a vaccine.

opposed to 29%, and consumers who do receive a penalty have a slightly smaller one, $43 as opposed to $52 for the average game. Figure 7 demonstrates that the model fits the observed WTP changes in the data.

Table 7: Estimated preference change parameters for the transition from 2019 to the vaccine state. Standard errors calculated using the bootstrap.

Parameter	Value	Std. Err
ρ_1^{2019}	0.60	0.02
ρ_2^{2019}	43.20	4.58

E Additional Robustness Checks and Counterfactuals

Resale Market Efficiency

This subsection considers how the efficiency of the resale market in the model compares to the effiency of the observed resale market. The issue is worthwhile because there are sources of inefficiency in the real world—such as dynamics, imperfect browsing, and potentially market power—that are not explicitly modeled. In contrast, the resale market in the model is static and, by Assumption 3, has a single clearing price.

A concern is whether the modeled resale market is too efficient, and whether it affects the comparison to refunds.

There are two important items to note before presenting robustness checks. First, the resale market in the model has a source of inefficiency, the resale friction s_{ij}, which also creates dispersion in the effective prices paid by resale buyers. For that reason, the clearing price assumption should not be treated as an indicator that the resale market is perfectly efficient.

Second, I do not have a full model of the resale market that would quantify its inefficiencies. Although I can measure inefficiency in the model, there is no benchmark indicating the correct amount of inefficiency.

Welfare Gains from Resale. Given the lack of a benchmark resale market, one useful exercise is to compare the welfare gains from resale in a market with and without the resale frictions s_{ij}. Essentially, if the resale market were truly efficient and lacked price dispersion, how would the results differ from those in the model?

In an out-of-equilibrium comparison, I take the first-period allocation from estimation as given and compare welfare created through resale to the welfare that would be created if there were no resale frictions. Welfare gains with estimated frictions are only half as high as they would be without frictions. Total welfare generated in the second period would be 11.4% higher without frictions.

The counterfactual experiment with λ_s set to zero provides an equilibrium comparison. Total welfare increases by 1.6%, again based on an idiosyncratic shock rate of 8%. The difference is large enough for resale (without resale frictions) to become more efficient than refunds.

Price Dispersion. A second exercise is to compare observed price dispersion to the distribution of effective prices paid in the resale market, defined as $p_j^r + s_{ij}$. (Dispersion does not depend on seat quality; I compare dispersion for the common value V realized in the data.) Price dispersion is not a perfect metric because a model can capture inefficiency without necessarily matching price dispersion, but it is easily observed and so provides a convenient benchmark.

Tables 8 and 9 show various percentiles of the distribution of resale prices minus the median (using effective resale prices paid by buyers for the model, which include incurred frictions).

In general, the model implies less price dispersion than is observed in the data, but the results depend on the game. The model-implied results vary more across games,

Percentile	Game 1	Game 2	Game 3	Game 4	Game 5
15	-3.5	-7.0	-3.5	-5.5	-2.5
25	-2.5	-6.0	-2.5	-4.5	-2.0
75	6.0	13.5	6.0	12.5	4.0
85	10.5	24.5	10.0	23.0	8.0

Table 8: Model-implied dispersion from the median in effective resale prices.

Percentile	Game 1	Game 2	Game 3	Game 4	Game 5
15	-11.5	-10.4	-11.3	-11.4	-8.6
25	-7.9	-5.9	-8.4	-7.5	-4.6
75	9.1	8.2	11.2	10.3	5.9
85	14.4	14.3	17.5	15.9	11.2

Table 9: Observed dispersion in resale transaction prices.

so popular games like 2 and 4 have similar or larger interquartile price spreads in the model than in the data. In all other cases, the price dispersion in the data is greater. The model matches a meaningful amount of the price variation observed in the resale market.

Partial Entry to the Resale Market

One significant assumption in the article is that all tickets are used. In reality, we expect that some tickets are neither used nor resold, and the number of wasted tickets affects the welfare gains from reallocation.

For the article's central comparison of resale and refunds, the salient issue is whether there will be more or less waste with one policy. In theory, the effect could go either way. Resale markets are more of a hassle than pushing a button that says "claim refund," so there could be more reallocation with refunds. But resale prices should generally be higher than the primary seller's partial refund (so that fewer consumers ask for refunds when the team is bad), leading to relatively more reallocation through resale. Which effect dominates is an empirical question that I cannot resolve with the data at hand.

Nonetheless, it is worth looking at simulations with partial take-up to gauge the magnitude of the effect. I introduce a probability that consumers forget to reallocate.

Assumption 4. *With probability ϕ, consumer i forgets to resell tickets or claim*

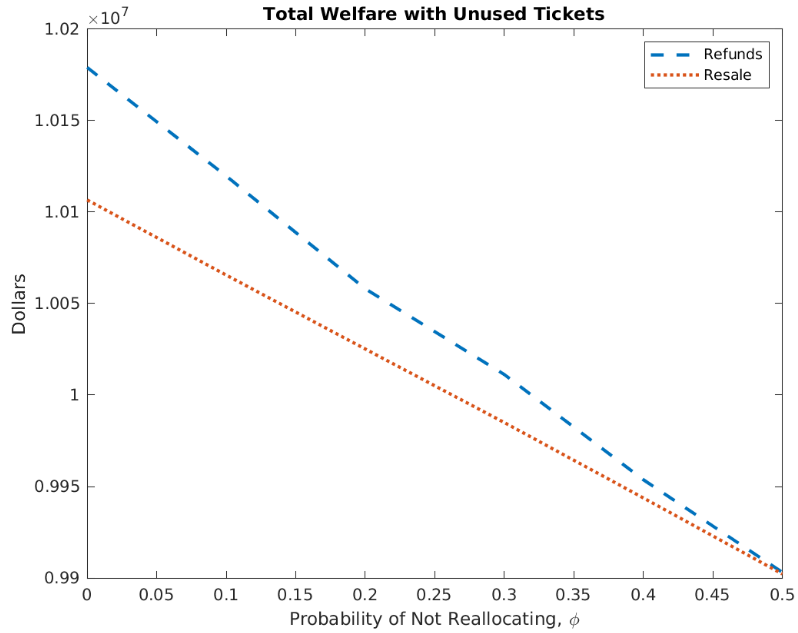

Figure 8: Total welfare in the resale and refunds counterfactuals as ϕ varies.

a refund. Consumers know ϕ in the first period of the model but do not learn its realization until the second period.

The probability ϕ is a simple way to capture the idea that consumers face hassle costs that prevent them from entering the resale market or from seeking a refund. Because the rate of idiosyncratic shocks ψ is fixed, the volume of reallocation decreases as ϕ increases.

To gauge the magnitude of the effect, I run the counterfactual experiments for values of ϕ ranging from 0.1 to 0.5, in increments of 0.1.

The results in Figure fig:partial-entry-cf confirm that the takeup rate for reallocation could affect which aftermarket policy maximizes welfare. For example, although refunds produce higher total welfare at each level of ϕ considered, society prefers resale with no waste to refunds with a probability of forgetting exceeding 0.15. In general, refunds needs a higher rate of waste to produce similar total welfare as refunds. The gap narrows as ϕ grows because the volume of idiosyncratic shocks ψ is fixed, causing welfare with both policies to approach the welfare level without reallocation as ϕ approaches one.

Unobserved Reallocation

There is undoubtedly reallocation outside of the primary and online resale markets, like when people give tickets to friends or sell offline. However, those transactions are unobserved and I have not been able to find data on their scope.

The issue is potentially important to the counterfactual results because the refund policy would shut down the alternative channels. (The primary seller must prohibit private transfers of tickets to prevent resale when implementing a refund.) If reallocation in the unobserved channels is efficient, then the counterfactuals understate the performance of resale. If it is not, then refunds should hold a bigger advantage. The issue also affects whether the model features the right degree of reallocation.

Unobserved reallocation thus warrants attention, but the lack of data raises challenging questions in designing a robustness exercise. The central issue, which is pivotal to the results, is whether the consumers who get tickets through alternative channels have high values. Beyond that, any exercise requires an assumption about the volume of reallocation in alternative channels. Because consumers are forward-looking when buying season tickets, it also matters whether the alternative channels are giveaways or transactions and, if transactions, at what price.

An additional issue is whether the consumers who give tickets to friends would bother to collect a refund. If they currently give away tickets for free, it is not certain they would request the primary seller's partial refund. The implications of partial takeup are addressed by the partial entry counterfactual above, so I do not consider them here.

The counterfactual requires a handful of new assumptions.

Assumption 5. *Consumers receive idiosyncratic shocks at the rate ψ', where $\psi' = 1.5\psi$. Of the consumers receiving idiosyncratic shocks, one-third reallocate through alternative channels.*

Assumption 6. *Consumers are randomly selected to reallocate in alternative channels. Specifically, among consumers who receive idiosyncratic shocks, one-third are randomly selected to reallocate in alternative channels. Consumers do not know if they will reallocate through alternative channels before the second period.*

In the absence of data on the volume of reallocation in alternative channels, Assumption 5 supposes that reallocation through alternative channels is half as large as the reallocation through resale in the data. Although the estimate is not precise, it shows how the market would operate if alternative channels are significant. Assumption 6 clarifies that consumers cannot anticipate in which market they will reallocate. The assumption is necessary to keep the model tractable. Without it, consumers who would use alternative channels might make different choices in the first period.

Assumption 7. *Consumers who do not purchase season tickets are randomly sampled to be buyers in the alternative channels. Only consumers whose values ν exceed some threshold $\bar{\nu}$ are eligible to be selected. Exactly as many consumers are sampled as are needed to transact with the fraction $\psi' - \psi$ of season ticket holders who reallocate through alternative channels.*

Assumption 7 captures the idea that consumers who are not interested would not be offered tickets. I carry out the counterfactual for several thresholds $\bar{\nu}$.

Assumption 8. *Buyers and sellers for game j and quality q transact in alternative channels at price* $\min\{\alpha_j(\bar{\nu} + \gamma_q + V), \frac{3}{4}p^r_{jq}(V)\}$.

Assumption 8 provides the transaction price that determines the division of surplus in alternative channels. It ensures that all participants have positive surplus in the transaction and matches the likely scenario where sales to friends or acquaintances are discounted. There are no fees or frictions so that the only parameter controlling the inefficiency of unobserved resale is the threshold $\bar{\nu}$. The lack of frictions is not without consequence, and boosts the performance of resale. The transaction price is necessary because it enters the season ticket decision for forward-looking consumers.

As before, consumers are forward-looking. It follows that the new assumptions affect decisions in the first period. When assessing values for season tickets, consumers know the probability of receiving an idiosyncratic shock. They also know that they would resell two-thirds of the time with a shock and would use the alternative resale channels one-third of the time. Similarly, when deciding whether to wait for tickets, consumers know the threshold $\bar{\nu}$ and consider the chance that they would be offered tickets.

Because of the many assumptions and unknowns, the modified model cannot assess how wrong the main counterfactual results are. But it provides information on two questions that help assess the main counterfactual results. First, how different are the main resale results from the likely results when alternative channels matter? And second, how would refunds compare to resale when both reflect reallocation through unobserved channels?

Before presenting the results, there are two weaknesses of the experiment worth discussing. First, consumers are randomly selected to participate in the alternative resale channels. It is not an optimal choice for sellers, and buyers may be conscripted even if another choice is preferable. Second, consumers do not know in advance if they are the type that would engage in resale through alternative channels, which is unlikely in reality.

Both choices are meant to simplify the necessary changes to the model. Rationalizing participation in alternative channels would require additional assumptions on how unobserved resale works, and would require significant new machinery in the model. Making participation in alternative channels a dimension of consumer types would make it difficult to control the number of consumers who participate in alternative channels, and would also require larger changes to the code base.

	Resale (Benchmark)	Refunds	$\bar{\nu} = 10$	$\bar{\nu} = 20$	$\bar{\nu} = 30$
Tot. Welfare (mn)	\$10.11	\$10.05	\$9.95	\$10.02	\$10.06
Profit (mn)	\$7.17	\$7.23	\$6.86	\$6.91	\$6.95
Cons. Welfare (mn)	\$2.84	\$2.82	\$2.99	\$3.00	\$3.01
Resale Fees (mn)	\$0.10	\$0.00	\$0.10	\$0.10	\$0.10

Table 10: Counterfactual results for the main resale experiment, a partial refund experiment with the rate of idiosyncratic shocks used in the unobserved resale experiments, and unobserved resale experiments with varying thresholds for consumers to randomly receive tickets.

Results. Results of the counterfactual are shown in Table 10. The leftmost column provides the main resale results from Section 7 for comparison. The refunds column implements the same partial refund scheme from Section 7 with the new idiosyncratic shock rate ψ' to provide a fair comparison for the counterfactuals with alternative channels. The remaining columns consist of the counterfactuals of interest and differ by the value threshold $\bar{\nu}$ used.

The first comparison is between the resale results with different values of $\bar{\nu}$. As the value threshold climbs from 10 (a consumer without particular fondness for college football) to 30 (a consumer who is on the fence about buying season tickets), total welfare and profit change significantly relative to the scope of reallocation. The details of offline reallocation are therefore important.

The next comparison is between the resale results and the modified refund counterfactual that reflects the higher rate of idiosyncratic shocks. Refunds are clearly more profitable and worse for consumers than resale, but whether they are more efficient depends on the true value of $\bar{\nu}$. However, given that resale only produces a slightly higher level of total welfare when offline resale is limited to high-value consumers, it is likely that refunds remain more efficient.

The final comparison is between the resale results in the main text and resale with offline reallocation. If offline reallocation is significant, the resale results in the main text would exaggerate the profit and welfare earned through resale, although consumer welfare would be understated.

The overall conclusion from the exercise is that offline resale, if significant, would change the levels of predicted welfare and profit in the market, but that it is not likely to change the assessment that refunds are more efficient and profitable.

Integrated Resale

In recent years, there has been understandable interest in integrated resale, defined as offering both primary and resale tickets on a single platform. The most notable example is Ticketmaster, which controls primary market sales for a variety of high-profile events and now shows "Verified Resale Tickets" alongside primary market inventory on its platform.

Integrated resale could have a range of effects. One is to enhance Ticketmaster's market power, enabling it to charge higher fees. Ticketmaster does not disclose its fee structure, but adds 23% of the before-fee price to what buyers pay on a listing for an upcoming event at the time of writing. If Ticketmaster charges resellers a similar amount as StubHub, its fees are indeed higher. A second effect is that, by making resale tickets easier to find, integration may reduces the frictions associated with participating in the resale markets. A third effect is that any of the other changes may affect brokers' incentives to participate.

The empirical model of resale in this article can be modified to feature an integrated resale market, but it cannot provide evidence on the effects above. Brokers are not prominent in the data and so do not feature in the model. The model is thus not tailored to determine how integrated resale would affect their behavior. For frictions, the model allows me to assess the effect of different levels of frictions, but there is no basis for judging how resale frictions should change in an integrated resale market. As in other robustness checks, I provide a range of estimates without declaring which is correct.

The model also cannot assess how an integrated seller would change fees. One reason is that I do not model competition between resale platforms, which could affect fees. The other is that the model cannot find the optimal fee, even if the integrated seller were a monopolist in the resale market, because there is no intensive margin to reallocation. Without an intensive margin, a change in fees can be exactly offset by a change in the season ticket price in the model, resulting in the same allocation and profit.

In more detail, consumers only resell in equilibrium when they receive an idiosyncratic shock. When buying season tickets of quality q, they thus expect to receive $\psi(1-\tau)\alpha_j(\bar{p}_j^r+\gamma_q)$ in resale revenue for each game j. (For convenience, let $\bar{p}_j^r$ be an average over all realizations of V.) When deciding whether to buy season tickets, consumers consider the price p_{Bq}, defined in equation (12), minus resale revenue. The difference is

$$\begin{aligned}
&p_{Bq}-\left(\sum_j \alpha_j\right)\psi(1-\tau)(\bar{p}_j^r+\gamma_q)\\
=&\left(\sum_j \alpha_j\right)\big(p_B+(1-\psi\tau\gamma_q)-\psi(1-\tau)(\bar{p}_j^r+\gamma_q)\big)\\
=&\left(\sum_j \alpha_j\right)\big(p_B+(1-\psi)\gamma_q-\psi(1-\tau)\bar{p}_j^r\big)\,.
\end{aligned}$$

An increase in fees of $\Delta\tau$ can thus be offset exactly by a reduction in p_B of $\psi\Delta\tau\bar{p}_j^r$, which would result in the same season ticket choices, resale prices, allocations, and profit for an integrated reseller. For this reason, eliminating the fee τ only moves revenue from the resale market to the primary market.

As before, a counterfactual still shows the range of possible changes if the resale friction changes. The specific assumptions are as follows.

Assumption 9. *The profit-maximizing primary market seller controls the entire re-*

sale market and earns all fees. Fees are fixed at the level τ observed in the data.

Although fees now contribute to primary market profit, they are still listed separately in the counterfactual results. The assumption that the primary market seller controls the entire resale market ignores the fact that there may be competing platforms.

The assumption that fees are fixed at the level observed in the data is primarily motivated by the fact that the model cannot identify an optimal fee. However, the integrated seller's choice of fee could also be constrained if it competed with other resale platforms.

	Not Int.	Int.	Int., $\lambda_s = 60$	Int., $\lambda_s = 40$	Int., $\lambda_s = 20$
Total Welfare (mn)	\$10.11	\$10.11	\$10.12	\$10.14	\$10.18
Profit (mn)	\$7.17	\$7.27	\$7.28	\$7.30	\$7.35
Consumer Welfare (mn)	\$2.84	\$2.84	\$2.84	\$2.84	\$2.84
Resale Fees (mn)	\$0.10	\$0.10	\$0.10	\$0.11	\$0.12
Season Ticket Base Price	\$31.82	\$31.83	\$31.89	\$32.08	\$32.34
Single Game Base Price	\$42.22	\$42.22	\$42.17	\$41.78	\$39.85

Table 11: Counterfactual results with an integrated resale market and the fee τ fixed at the observed level in the data.

The results of the exercise are shown in Table 11. The leftmost column shows the counterfactual results for resale without an integrated resale market. The next column shows results for an integrated resale market with the value of resale frictions observed in the data, and subsequent columns differ by varying the level of frictions.

The results show that, as in the counterfactual in Section 7 where the fee τ is removed, integrated resale is a pure transfer of surplus to the primary market seller. In the main comparison, between the leftmost columns, the only difference is that the primary market seller's profit is higher by the amount of fees. Therefore, in the empirical model, integrating the resale market does not change the primary market seller's incentives enough to affect the allocation of tickets. The results in the rightmost three columns are driven by the reduction in frictions, which raise total welfare. The increase in welfare is captured entirely by the primary market seller.

Refund Levels

In this section, I evaluate the partial refunds counterfactual with different levels of the refund.

In the robustness exercise, the primary market seller sets optimal prices while taking a level of the refund as given. The approach differs from that used in Section 7, where the refund r_j was set as 30% of the season ticket price for zone 5 tickets to game j. Thus, in the main text, the primary seller's season ticket prices and the level of the refund were linked. The link was not consequential because the level of the refund was relatively low.

In contrast, the robustness exercise considers cases where the size of the refund makes the link between the refund level and prices strategically important. For example, if the refund were set at 90% of the season ticket price to game j, the primary market seller will change its prices to affect the level of the refund.[4] To remove distortions and focus on the level of the refund, I take the level of the refund as given.

Recall that the main counterfactual took the refund to be 30% of the optimal per-game season ticket price for zone 5 seats. The refunds in this exercise use the optimal per-game season ticket prices as a benchmark. The primary market seller offers refunds ranging between 10% and 90% of those prices. The 30% refund produces the same results as in the main text. Results are in Table 12.

	10% Refund	30% Refund	50% Refund	70% Refund	90% Refund
Total Welfare (mn)	\$10.17	\$10.18	\$10.18	\$10.18	\$9.68
Profit (mn)	\$7.33	\$7.34	\$7.34	\$7.41	\$7.45
Consumer Welfare (mn)	\$2.84	\$2.84	\$2.84	\$2.77	\$2.22
Season Ticket Buyers (1000)	26.87	27.03	27.18	27.49	28.24
Avg. Returned Tickets	0.00	0.00	2.00	100.00	1490.00
Season Ticket Base Price	\$30.21	\$30.71	\$31.19	\$32.06	\$33.79
Single Game Base Price	\$41.07	\$40.98	\$41.37	\$41.37	\$42.07

Table 12: Results of the partial refunds counterfactual with varying levels of the refund. Refund levels are percentages of the optimal primary market price in the baseline partial refunds counterfactual (where the refund level is set at 30%). Average returned tickets denotes the average number of consumers who request refunds to each game because they have low values rather than schedule conflicts.

The most notable finding is that the results are robust when the level of the

[4] Its incentives to distort prices would also vary based on whether the refund is proportional to season ticket or single-game prices.

refund is low. Setting the refund as high as 50% would have virtually no effect on the main outcomes of interest, and would cause virtually no consumers to return tickets without receiving an idiosyncratic shock.

The picture starts to change when the refund climbs past 70%. The most notable effect is that the primary market seller sells more season tickets and manages to earn higher profit, although consumer and total welfare fall. The increase in profit is counterintuitive: high refunds should lead to many redemptions in low-demand states, which would lower both total surplus and profit.

How does profit increase? The likely answer is that the refund facilitates price discrimination. When the refund increases, season ticket buyers with relatively low values are likely to return tickets in states with a low realization of V, but those with high values are unlikely to request any refunds. Consumers with modest values thus have a willingness to pay for season tickets based on their values when the team is reasonably good—they will just request refunds if the team underperforms. They are willing to pay more for season tickets than they would with a lower refund. The same is not true for high-value consumers who would not request refunds in low-demand states.

The primary market seller does not profit directly from the higher willingness to pay of consumers with modest values. (The increase in willingness to pay reflects higher refunds that the primary seller will have to pay.) However, the primary market seller benefits indirectly because marginal buyers of season tickets are willing to pay more, letting it charge a higher price without losing as many buyers. The strategy leads to losses from consumers who request refunds, but the losses are outweighed by the increase in revenue from inframarginal buyers of season tickets who pay higher prices. The result is higher season ticket prices, lower total welfare, and higher profit.

There are reasons to be cautious when interpreting the increase in profit at high levels of the refund. First, whether the strategy is reasonable depends on how often consumers would request refunds, which the data cannot show. If there are frictions associated with requesting refunds, then it may not be as effective in separating high- and low-value buyers of season tickets as the model suggests. Second, I am not aware of any examples of a seller using such a strategy in ticket markets.

Screening Counterfactuals

The discussion of the resale counterfactual in Appendix A noted that some type L consumers purchased season tickets when δ_L was set to its upper bound. As a result, there may be different counterfactual results for different identified sets of parameters. In this section, I show that the results of the counterfactual are virtually the same at lower values of δ_L.

I run the counterfactual experiment again after setting δ_L to -1000. At the chosen value, no consumers of type L purchase season tickets. It is therefore a useful comparison to assess the spread of outcomes within the identified set of parameters. The results are shown in Table 13 and hardly differ. The partial identification of δ_L therefore has a negligible effect on interpreting the counterfactual results.

	δ_L Upper Bound	$\delta_L = -1000$
Profit (mn)	\$6.76	\$6.77
Consumer Welfare (mn)	\$2.67	\$2.66
Total Welfare (mn)	\$9.44	\$9.44
Resale Fees (mn)	\$0.01	\$0.01
Non-Refund. S. Tix (1000)	25.87	25.83
Type L S. Tix (1000)	0.24	0.00

Table 13: Counterfactual results for the resale counterfactual in the screening application when δ_L is set to the upper bound of the identified set (-203) and a value low enough that no type L consumers purchase season tickets ($-1,000$).

F Survey

I surveyed 250 Americans under the age of 50 and 250 Americans aged 50 or over, ultimately receiving a total of 457 usable responses. I distributed the survey through Prolific.co, an online survey distribution platform. Respondents were paid \$9.34 per hour and live in nine states that each have one dominant college football team: Arkansas, Georgia, Louisiana, Michigan, Minnesota, Nebraska, Ohio, West Virginia, and Wisconsin. Respondents from each state were asked to consider one ticket for that team throughout the survey.

I asked for the amount they are willing and able to pay in four scenarios: (i) the 2019 season, (ii) a covid-19 vaccine, (iii) no vaccine but the number of cases falls

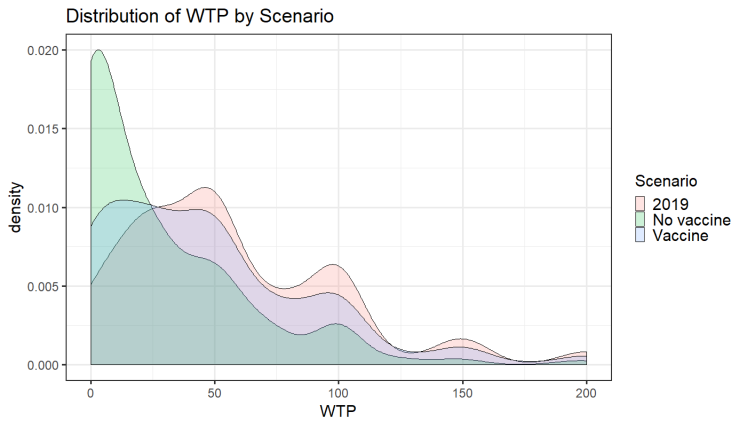

Figure 9: Distribution of reported willingness to pay without social distancing in 2019, with a vaccine, and with no vaccine.

below the CDC's near-zero benchmark, and (iv) no vaccine and the number of cases is above the CDC's near-zero benchmark.

The CDC's benchmark for a near-zero number of new cases is 0.7 new cases per 100,000 people. Respondents were given the benchmark and a practical illustration, that a 25,000-seat stadium filled with randomly selected people would contain an average of 2.5 sick people if each case lasts two weeks. They were also told that the true number of infected people would be lower, on average, because some people would know they are ill and decide not to attend.

The survey includes respondents with a wide range of reported WTP. Figure 9 shows the distribution of reported WTP for three scenarios without social distancing: a 2019 baseline, a state with a vaccine, and a state without one. In each state, some consumers report values for tickets exceeding $50 and $100.

In the absence of a true measure of the probability of each scenario in the future, I ask respondents how likely they consider each one at three future dates. The average percent chances are shown in Figure 10. Respondents do not expect a vaccine in January 2021, but think the chance exceeds 40% in September 2021 and 60% in January 2022.

Figure 11 shows that the distribution of reported WTP is similar for the near-zero and above near-zero scenarios.[5] The distributions are not exactly the same—consumers are more reluctant to attend when there are more cases—but the differ-

[5]The figure shows reported WTP without social distancing. The analogous figure with social distancing is similar.

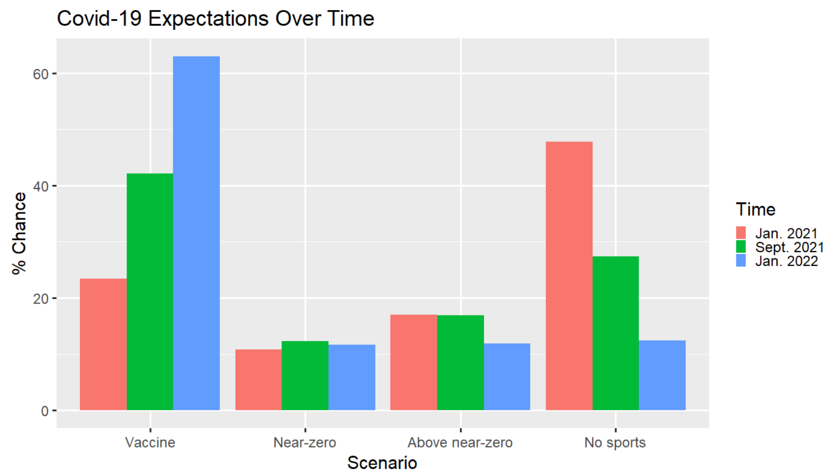

Figure 10: Average reported percent chance of each scenario occurring in each month.

ences are small enough for the two to be consolidated into a single state without a vaccine. I consolidate WTP as a weighted average, taking the relative probability of the states in September 2021 as the weights.

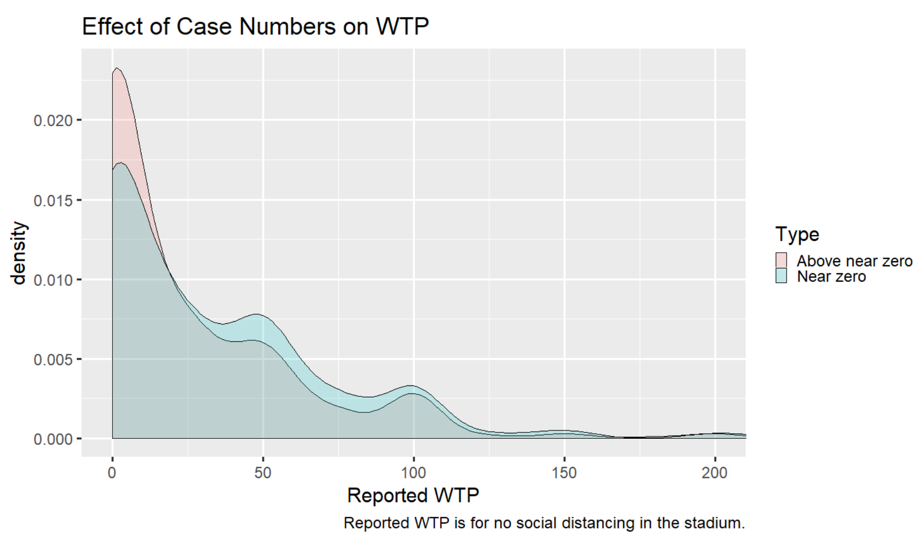

Figure 11: WTP distributions with near-zero and above near-zero levels of cases.

Figure 12 shows that the distribution of reported WTP is also similar with and without social distancing. As before, there are some changes, but they are not large enough to treat separately. I use reported values without social distancing because distancing would greatly reduce the number of tickets the primary market seller can offer.

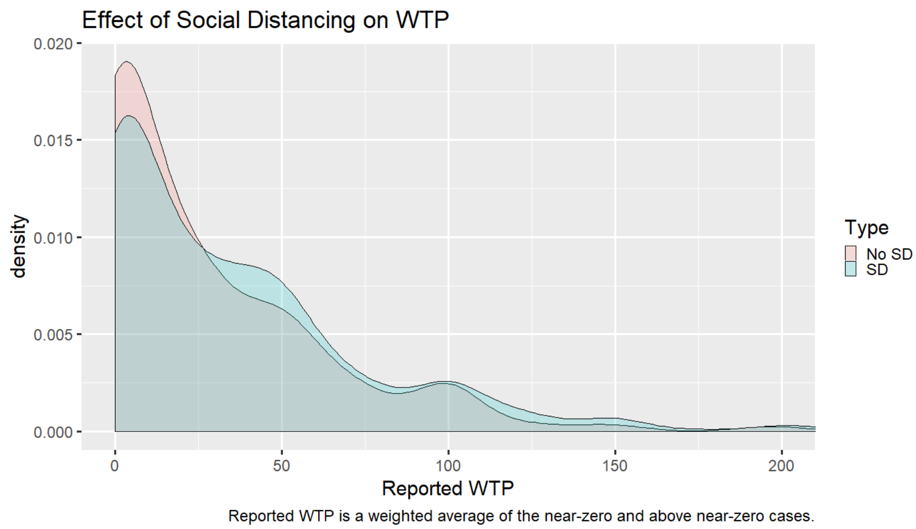

Figure 12: WTP distributions with near-zero and above near-zero levels of cases.

Surprisingly, demographics were not an important determinant of the change in WTP across states. I evaluated regression models of the form

$$\Delta WTP_i = \alpha Age_i + \beta Race_i + \gamma State_i + \varepsilon_i, \tag{14}$$

where Age_i is a set of age dummies (with decade-long bins, e.g. ages $30--39$), $Race_i$ is a set of dummies for race, and $State_i$ is a dummy for the state of the respondent. The response variable is measured both as an absolute number of dollars and as a percentage of initial WTP. Lower values denote greater sensitivity to the state without a vaccine. Results are shown in Table 14.

Although all groups were more sensitive than the reference group of respondents aged 22–29, those aged 50 and over did not report greater sensitivity to the state without a vaccine than those aged 30–49. (Respondents 70–79 are not numerous and only show a stronger response in one model.) Responses vary by race, but no coefficients are significant and the groups with large changes have few respondents. Because the covariance of value differences with demographics is not a critical feature of the data, I make value changes independent in the empirical model.

The full survey is included below.

Table 14: Regression output for equation (14).

	Dependent variable:	
	Value DIfference	Value Difference (%)
	(1)	(2)
Age 30-39	−12.821	−0.220
	(8.713)	(0.088)
Age 40-49	−12.419	−0.051
	(8.806)	(0.089)
Age 50-59	−19.863	−0.169
	(8.911)	(0.090)
Age 60-69	−7.068	−0.125
	(9.574)	(0.097)
Age 70-79	−10.209	−0.390
	(10.736)	(0.109)
Asian	28.256	−0.156
	(33.078)	(0.335)
African American	43.199	−0.267
	(32.337)	(0.328)
Other	35.790	−0.495
	(36.947)	(0.374)
White	50.867	−0.034
	(30.960)	(0.314)
White, Asian	−36.666	−0.078
	(60.378)	(0.612)
White, African American	74.146	0.447
	(48.460)	(0.491)
Constant	−59.122	−0.213
	(32.724)	(0.331)
State Fixed Effects	Yes	Yes
Observations	382	382
R^2	0.070	0.114
Adjusted R^2	0.013	0.060
Residual Std. Error (df = 359)	41.706	0.422
F Statistic (df = 22; 359)	1.233	2.105^{***}
Note:	Age coefficients relative to respondents aged 22–29. Race coefficients relative to respondents who are American Indians or Alaska Natives.	

Event Expectations (General)

Start of Block: Intro

Q1 This study is conducted by Drew Vollmer, a doctoral student researcher, and his advisor, Dr. Allan Collard-Wexler, a faculty researcher at Duke University.

The purpose of the research is to design sales strategies that cope with uncertainty over the covid-19 pandemic. You will be asked about how much you would pay for tickets to an outdoor college football game under several scenarios related to covid-19. The survey should take 5-10 minutes.

We do not ask for your name or any other information that might identify you. Although collected data may be made public or used for future research purposes, your identity will always remain confidential.

Your participation in this research study is voluntary. You may withdraw at any time and you may choose not to answer any question. You will not be compensated for participating.

If you have any questions about this study, please contact Drew Vollmer. For questions about your rights as a participant contact the Duke Campus Institutional Review Board at campusirb@duke.edu.

End of Block: Intro

Start of Block: Block 4

Q16 In which state do you currently reside?

▼ Alabama (1) ... I do not reside in the United States (53)

End of Block: Block 4

Start of Block: WTP

JS

Q2
In this section of the survey, you will be asked how much you are **willing and able to pay for one ticket to a football game**. Your responses should be dollar amounts.

In some questions, you will be given a scenario related to COVID-19. You should respond with the amount you would pay if that scenario occurs. You should not consider how likely the scenario is.

Q3 What is the **maximum** you would be **willing and able to pay** for **one** ticket...

	Amount (dollars) (1)
one year ago, in Fall 2019? (1)	
if there had not been a global COVID-19 outbreak and the virus had not spread to the US? (2)	
if there is a widely available COVID-19 vaccine? (3)	

Q4
In the next two questions, suppose that there is **no COVID-19 vaccine**, but that fans are allowed to attend sporting events.

You will be asked to consider two levels of risk from the virus:

The CDC says that new cases are **near zero**. The CDC says that new cases are **more than near zero,** but **risk is low enough** to allow fans at sports games.

The CDC standard for new cases to be near zero is 0.7 new cases per 100,000 people or fewer. This means that filling a 25,000-seat stadium with randomly selected people would imply an average of **2.5 sick people** in the stadium if each case lasts two weeks. The true number of infected people at any event, however, would be lower because some people would know they are sick and would not attend.

---

*

Q5
Suppose that there is **no social distancing in the stadium**.

What is the **maximum** you would be **willing and able to pay** for **one** ticket if...

	Amount (dollars) (1)
the CDC says that the number of new cases is **near zero**? (4)	
the CDC says that the number of new cases is **higher than near-zero**, but that the risk from attending mass gatherings is **low enough** to allow fans at sports games? (5)	

---

*

Q6
Suppose that there is **social distancing in the stadium**.

What is the **maximum** you would be **willing and able to pay** for **one** ticket if...

	Amount (dollars) (1)
the CDC says that the number of new cases is **near zero**? (4)	
the CDC says that the number of new cases is **higher than near-zero**, but that the risk from attending mass gatherings is **low enough** to allow fans at sports games? (5)	

Q7

Suppose that fans can return their tickets if the number of new virus cases is higher than near-zero. Tickets are sold out, but there is a **wait list** in case fans who bought tickets return them because of the virus.

What is the maximum you would be willing to pay for a ticket on the wait list?

	Amount (dollars) (1)
No social distancing in the stadium (1)	
Social distancing in the stadium (3)	

End of Block: WTP

Start of Block: Probabilities

Q8
In this section, you will be asked about the likelihood of COVID-19 scenarios. Your answers should be *percent chances*. So, if you believe an outcome has a one-in-four chance of occurring, the percent chance is 25%.

JS *

Q34 What is the *percent chance* of each outcome in **January 2021**? Chances must sum to 100.

Current total: 0 / 100
________ There is a widely available COVID-19 vaccine. (1)
________ There is no COVID-19 vaccine and new cases are **near zero**, as defined by the CDC. (2)
________ There is no COVID-19 vaccine and new cases are **higher than near-zero**, but the CDC considers the risk from mass gatherings is **low enough** to allow fans at sports games. (3)
________ There is no COVID-19 vaccine, new cases are **higher than near-zero**, and the CDC judges that the risk from mass gatherings is **high enough** that fans cannot attend sports games. (4)

JS *

Q36 What is the *percent chance* of each outcome in **September 2021**? Chances must sum to 100.

Current total: 0 / 100
________ There is a widely available COVID-19 vaccine. (1)
________ There is no COVID-19 vaccine and new cases are **near zero**, as defined by the CDC. (2)
________ There is no COVID-19 vaccine and new cases are **higher than near-zero**, but the CDC considers the risk from mass gatherings is **low enough** to allow fans at sports games. (3)
________ There is no COVID-19 vaccine, new cases are **higher than near-zero**, and the CDC judges that the risk from mass gatherings is **high enough** that fans cannot attend sports games. (4)

Q35 What is the *percent chance* of each outcome in **January 2022**? Chances must sum to 100.

Current total: 0 / 100
________ There is a widely available COVID-19 vaccine. (1)
________ There is no COVID-19 vaccine and new cases are **near zero**, as defined by the CDC. (2)
________ There is no COVID-19 vaccine and new cases are **higher than near-zero**, but the CDC considers the risk from mass gatherings is **low enough** to allow fans at sports games. (3)
________ There is no COVID-19 vaccine, new cases are **higher than near-zero**, and the CDC judges that the risk from mass gatherings is **high enough** that fans cannot attend sports games. (4)

End of Block: Probabilities

Start of Block: Demographics

Q12 What is your year of birth?

____

Q13 What is your gender?

- ○ Male (1)
- ○ Female (2)
- ○ Prefer not to answer (3)

Q14 What is your ethnicity?

- ○ Hispanic or Latino/Latina (1)
- ○ Not Hispanic or Latino/Latina (2)

Q15 What is your race?

- ▢ White (1)
- ▢ Black or African American (2)
- ▢ American Indian or Alaska Native (3)
- ▢ Asian (4)
- ▢ Native Hawaiian or Pacific Islander (5)
- ▢ Other (6) ____

End of Block: Demographics