Estimation of inflow uncertainties in laminar hypersonic double-cone experiments

J. Ray, S. Kieweg, D. Dinzl, B. Carnes, V. G. Weirs, B. Freno, M. Howard, T. Smith, I. Nompelis and G. V. Candler

Online Supplemetary Material

In this document, we provide ancilliary information that supports the conclusions of the paper.

1 The Quality of Run 35's Inference

Here we check the quality of the estimated solution/PDF for Run 35. At the very least, the PDF should be able to reproduce the experimental data that was used to generate it. We take 100 random samples from the three-dimensional posterior PDF over (ρ_∞, U_∞, T_∞) and simulate flow over the double-cone for these realizations of the freestream conditions using our statistical emulators. These predictions are called the "push-forward posterior" (PFP) and, due to the use of emulators, are limited to the region of the fore-cone upstream the separation point. We repeat the same exercise using samples taken from the prior distribution (± 15 % variation about the freestream conditions in Table 1), with the aim of verifying whether the assimilation of observational data improved predictions over what was possible with the prior. In Fig. 1 (top row), we plot the predictions of the pressure and heat flux at the sensors due to samples from the prior and the posterior PDFs. The region between the first and third quartile (the inter-quartile range or IQR) is plotted in grey, whereas the 5^{th}- and 95^{th}-percentile predictions are plotted with dashed red lines. The solid blue line in the middle is the median prediction. Clearly, the PFP predictions are narrower, demonstrating the effectiveness of assimilating experimental data. A similar behavior is observed in the middle row, which plots the heat flux measured on the fore-cone. Finally, in the bottom row, we plot the Pitot-pressure predictions from the prior and posterior densities, normalized by the measured Pitot pressure. A value of unity implies a perfect match between the experiment and predictions. The bold line in the middle is the ratio of the median prediction to the experimental value. We see that the median prediction is over-predicted by about 5%. While the median predictions before and after calibration are not very different, the uncertainty in the prediction (vertical size of the box) is far smaller after calibration, i.e., in the PFP. In the bottom row (right), we plot the total enthalpy. Qualitatively, we see the same trend as for Pitot pressure, except the median predicted enthalpy is under-predicted by about 10%. The reason is as follows. As described

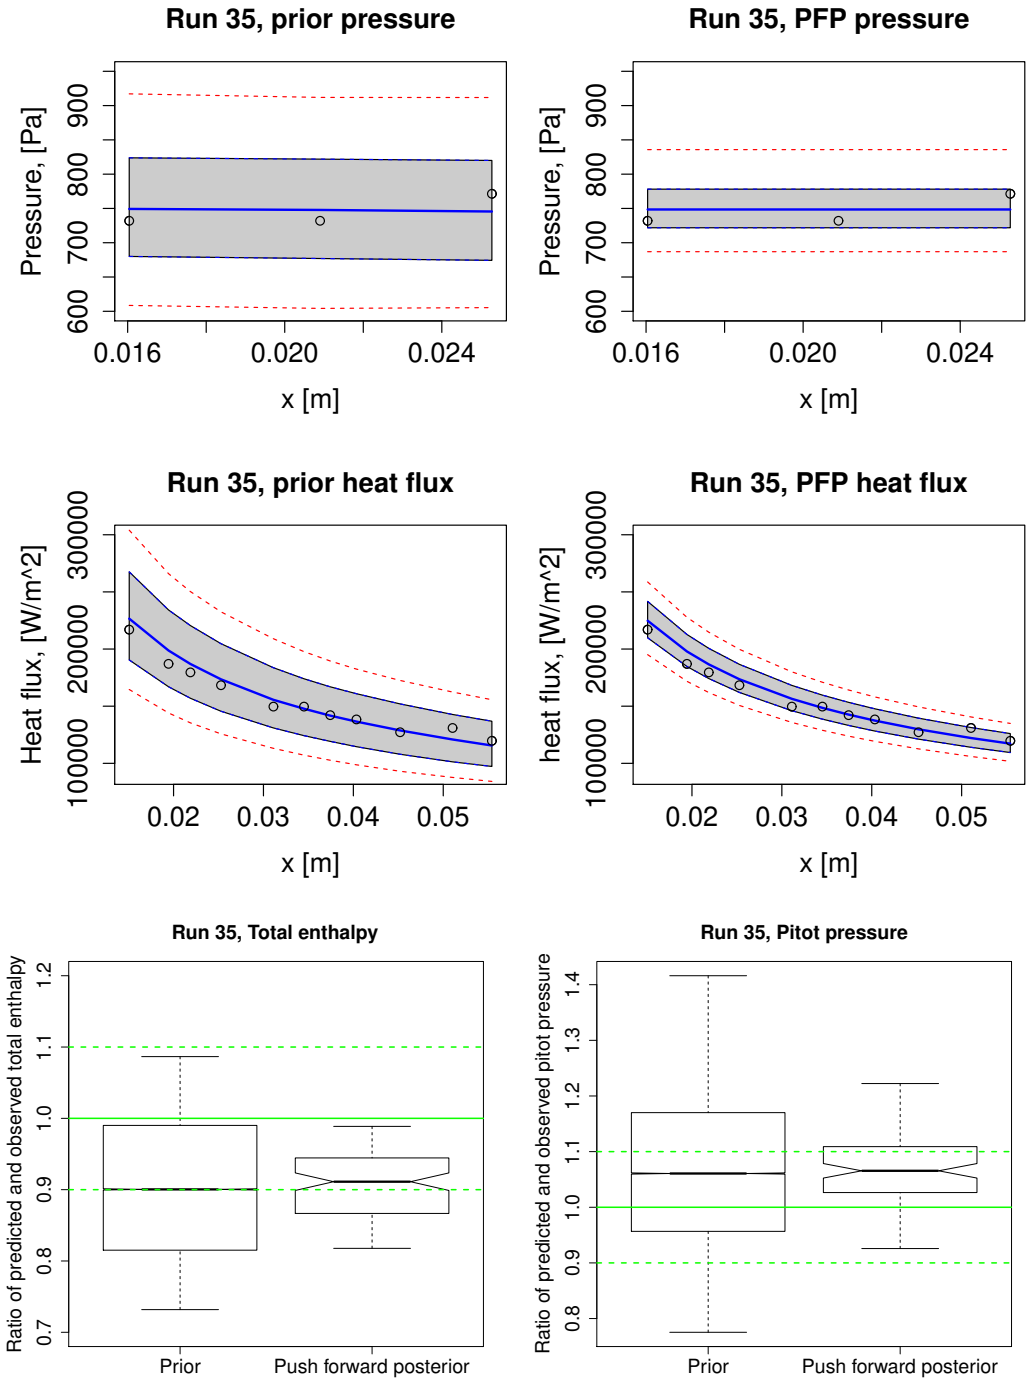

Figure 1: Top left: Predictions of pressure on the fore-cone using Θ samples drawn from the prior. The grey region is the IQR of predictions and the red dashed lines the 5^{th} and 95^{th} percentile predictions. The solid blue line is the median prediction. Top right: Pressure predictions from the PFP. Middle row: Same as the top row, but for heat flux on the fore cone. Bottom left: Box-and-whisker plot of predictions of Pitot pressure using Θ sampled from the prior and PFP and normalized by the measured value. Bottom right: Similar plots for h_0.

in Ref. [1], a part of the stagnation enthalpy in Run 35 was "frozen out" as vibrational energy upstream of the test section and is unavailable to the flow, i.e., the time-scale to equilibrate was too slow for the vibrational energy to contribute to the measured heat flux. Since our inference uses the heat flux to estimate the flow h_0, the "frozen out" fraction of the enthalpy could not be accounted for by the inference. Thus, the basic requirements of calibration, i.e., being able to estimate a tight posterior density for Θ seems to be met.

However, this does not automatically imply a high-quality estimation. This is deduced from the PFP plots for pressure and heat flux (Fig. 1). We see that after calibration, the IQR contains almost all the measurements, even though it is considerably narrower than its prior counterpart. However, since the IQR is the span between the 25^{th} and 75^{th} percentiles of the predictions, it should bracket about *half* the measurements, not *all* of them. Consequently, our calibration is *over-dispersed*. This is caused by the residual errors, called model-form errors, of the emulator. They are due to its approximate nature, errors in the experimental data, and the inherent shortcomings of the models in SPARC, including the boundary conditions (e.g., the assumption of uniform inflow) and the assumption of axisymmetry.

2 Validating against Run 6 and Run 7

To ensure that our performance with Run 35 was not a fortuitous result, we retest our framework with Run 6 and Run 7. The inflow conditions provided with the experimental dataset are in Table 1. We used pressure (7 sensors for both Run 6 and Run 7) and heat flux measurements (15 and 16 sensors respectively for Run 6 and Run 7) ahead of the separation point, along with the total enthalpy and Pitot pressures as listed in Table 1, to infer $(\rho_\infty, U_\infty, T_\infty)$ as a joint PDF. T_∞ does not influence the observables and consequently their marginalized posterior distribution was the same as the prior (a uniform distribution). Both the experiments are in "frozen" vibrational equilibrium and the total enthalpy is sufficiently low that N_2 can be treated as a perfect gas. The wall temperature was kept at 294 K and 297 K for Run 6 and Run 7 respectively.

In Fig. 2 we plot the marginalized PDFs for ρ_∞ and U_∞ for Run 6 (on top) and Run 7 (below). The MAP estimate, as well as the nominal freestream conditions stated in Table 1 are also plotted. The "nominal" freestream conditions we computed in Ref. [2] and assumed that the flow in the wind-tunnel test section was in thermochemical *non-equilibrium*. We see, as in Run 35, the ρ_∞ is not estimated well, but U_∞ is. The freestream conditions from Table 1 lie within the support of the PDF, and the difference between the MAP value, and those in Table 1 are within the measurement error bounds (also stated in Table 1). This is clearly seen in Table 2 where we tabulate the MAP values of the inferred freestream and compare with their experimental counterparts, including their uncertainties, obtained from

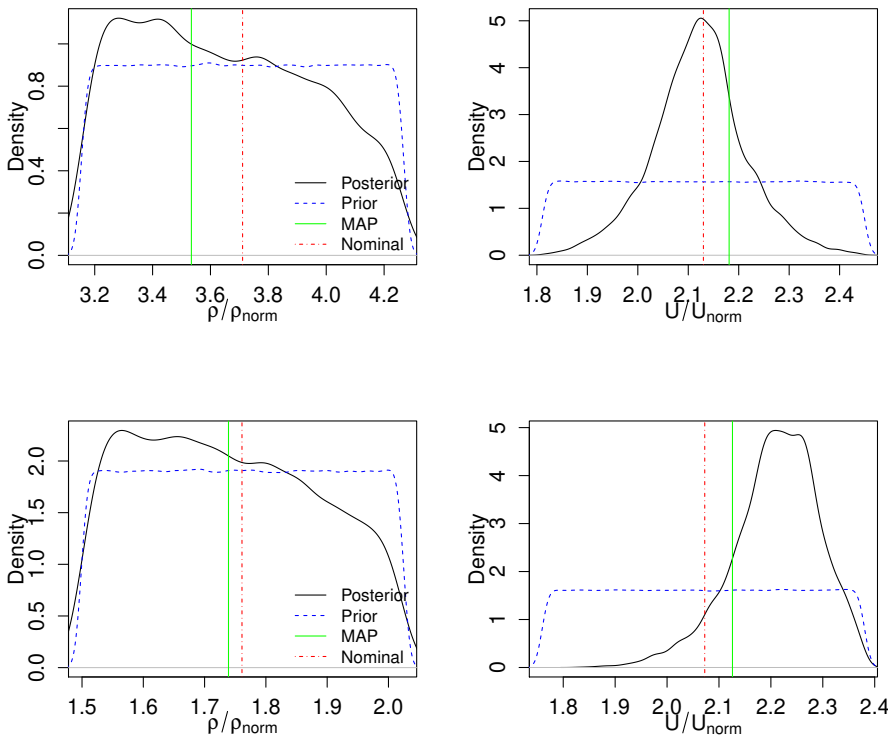

Figure 2: Top: Marginalized PDFs for ρ_∞/ρ_{norm} and U_∞/U_{norm} for Run 6. The MAP estimate and the nominal freestream conditions from Table 1 are also provided. Bottom: Marginalized PDFs for Run 7.

Table 1 .

Next we check the correctness of the PDFs for Run 6 and Run 7 developed in Fig. 2. We sample 100 $(\rho_\infty, U_\infty, T_\infty, T_{v_\infty})$ realizations from the PDF, simulate the double-cone experiment using SPARC and plot the ensemble of predictions (of pressure and heat flux) , along with the experimental values. These are in Fig. 3 (left). The 5^{th} and 95^{th} percentiles of the predictions with dotted lines, the IQR as the shaded region and the median prediction with a solid line and the prediction using the freestream conditions in Table 1 using a dashed line. The experimental measurements are plotted with symbols. We see that the extreme percentiles bound almost all the pressure measurements and all the heat flux ones. The box-and-whisker plot, at the right show the median predicted Pitot pressure and total enthalpy are slightly under-predicted (about 3%). In Fig. 3 (right), we perform the same check for Run 7. Both pressure and heat flux measurements are bounded by the extreme percentiles, and most fall within the IQR, except the measurements at the downstream end of the double-cone. We also see, from the box-and-whisker plots, that the Pitot pressure is not well predicted, but the total enthalpy is - the median prediction differs from the measured value (from Table 1) by 2%.

3 The Quality of Case 1's Inference

Next, we gauge the quality of the posterior PDFs for Case 1, plotted in Fig. 6 in the paper. In Fig. 4, we plot results from push-forward-prior and PFP predictions for pressure and heat flux on the fore-cone, as well as h_0 and P_{Pitot}. These predictions were performed using the statistical emulators seeded with 100 samples of Θ extracted from the posterior PDFs (plotted in Fig. 6 in the paper) and are limited to the fore-cone. The convention for plotting is the same as in Fig. 1. We see that the prior predictions are excessively wide (not a surprise given $\pm 15\%$ bounds) but seem to under-predict the heat flux and P_{Pitot} (see the median predictions and the tendency of the observations to cluster around the third quartile of the predictions). After Bayesian calibration, the spread in the predictions is much reduced, as is the bias between the median prediction and the observations (except for h_0). Thus, Bayesian calibration of Θ seems to have improved predictive skill of emulator-based simulations and could, potentially, do the same for SPARC. It also explains why we failed to bracket experimental data in Fig. 2 in the paper when using the "CUBRC priors" from Table 1 and succeeded in Fig. 4 when using the posterior PDF (from Fig. 6 in the paper). The support of the posterior distribution (in Fig. 6 in the paper) is wider, and spans a different range, than the support of the "CUBRC priors" (see Table 2).

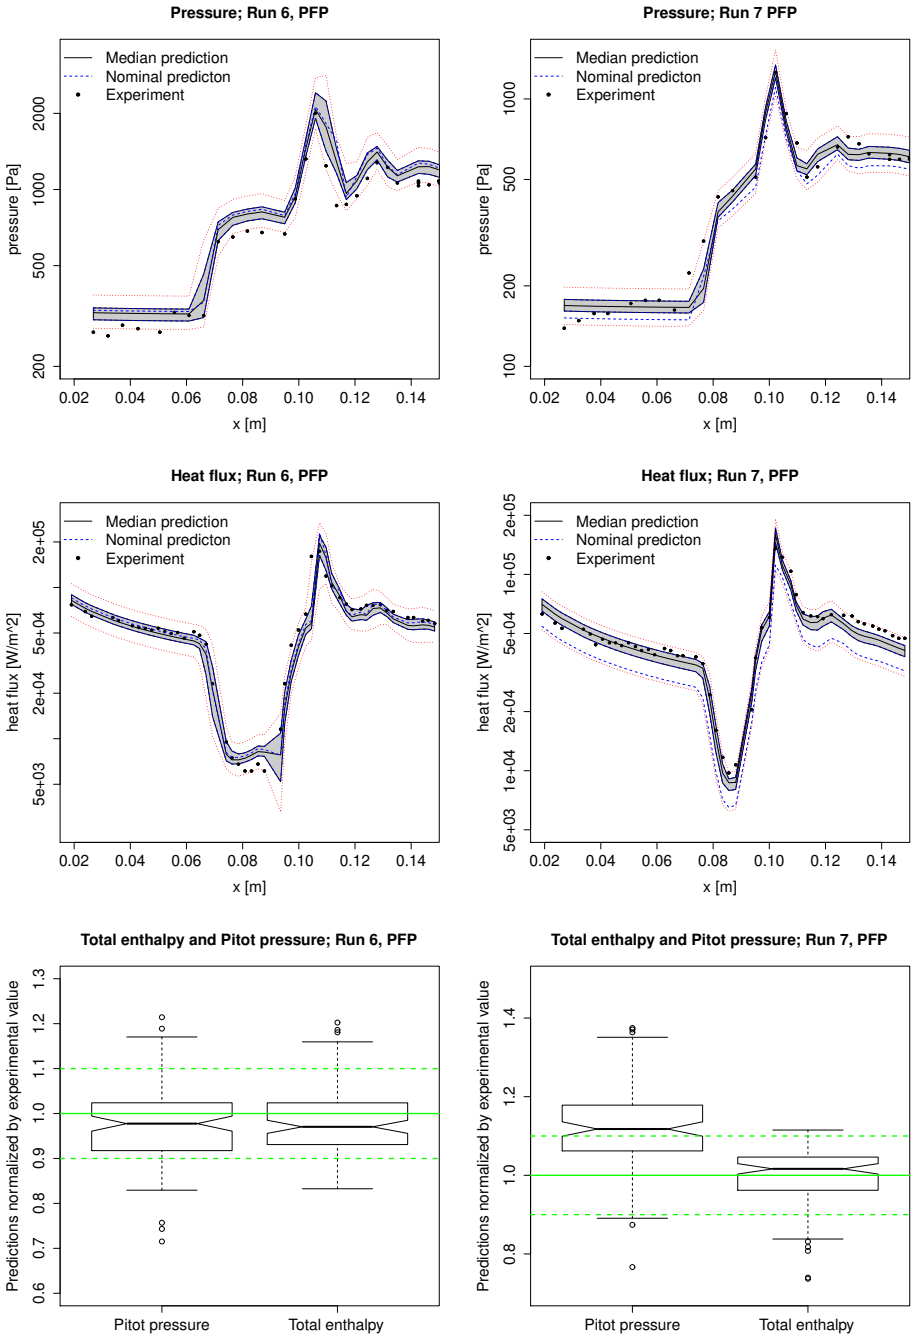

Figure 3: Left: PFP simulations for Run 6, using 100 samples of $(\rho_\infty, U_\infty, T_\infty)$ taken from the PDF in Fig. 2. We plot the pressure (top), heat flux (middle), total enthalpy and Pitot pressure (bottom). The 5^{th} and 95^{th} percentiles of the predictions with dotted lines, the IQR as the shaded region, the median prediction with a solid line and the prediction using the "nominal" freestream conditions using a dashed line. The experimental measurements are plotted with symbols. In the box-and-whisker plots, we plot the predicted enthalpy and Pitot pressure normalized by the measured values. The dashed lines are $\pm 10\%$ bounds. Right: The same, but for Run 7.

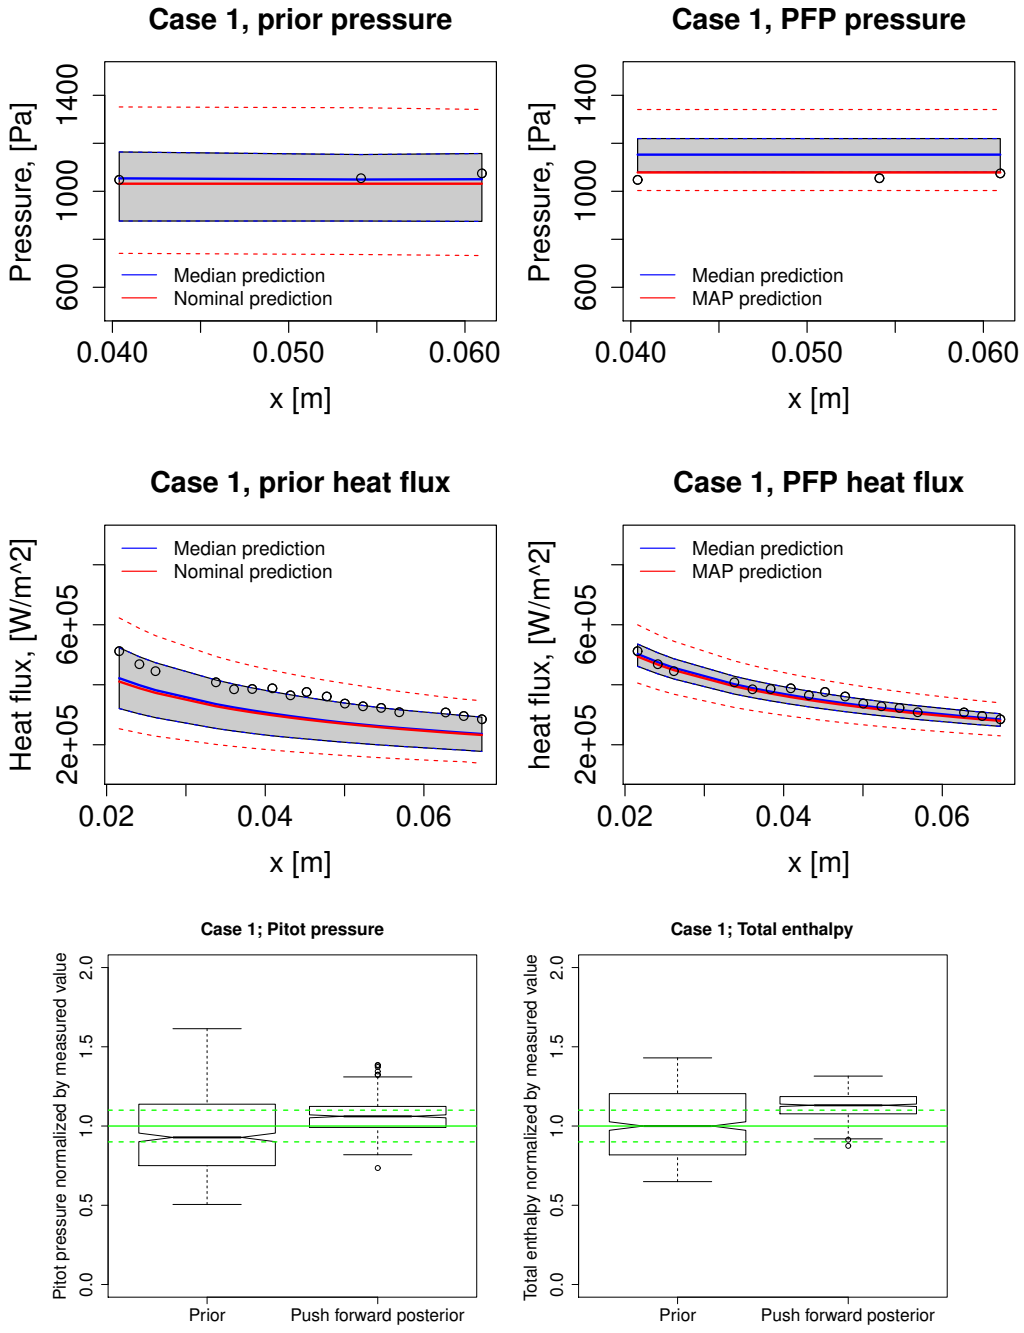

Figure 4: Top left: Predictions of pressure on the fore-cone using Θ samples drawn from the prior ($\pm 15\%$ bounds) for Case 1. The grey region is the prediction IQR and the red dotted lines the 5^{th}- and 95^{th}-percentile predictions. The solid blue line is the median prediction and solid red line the prediction using Θ_{MAP}. Top right: Pressure predictions from the PFP. Middle row: Same as the top row, but for heat flux. Bottom left: Box-and-whisker plot of predictions of Pitot pressure using Θ sampled from the prior and PFP and normalized by the measured value. Bottom right: Similar plots for h_0. The solid green horizontal line denotes a perfect match between predictions and observations. The dotted green horizontal lines are the $\pm 10\%$ limits.

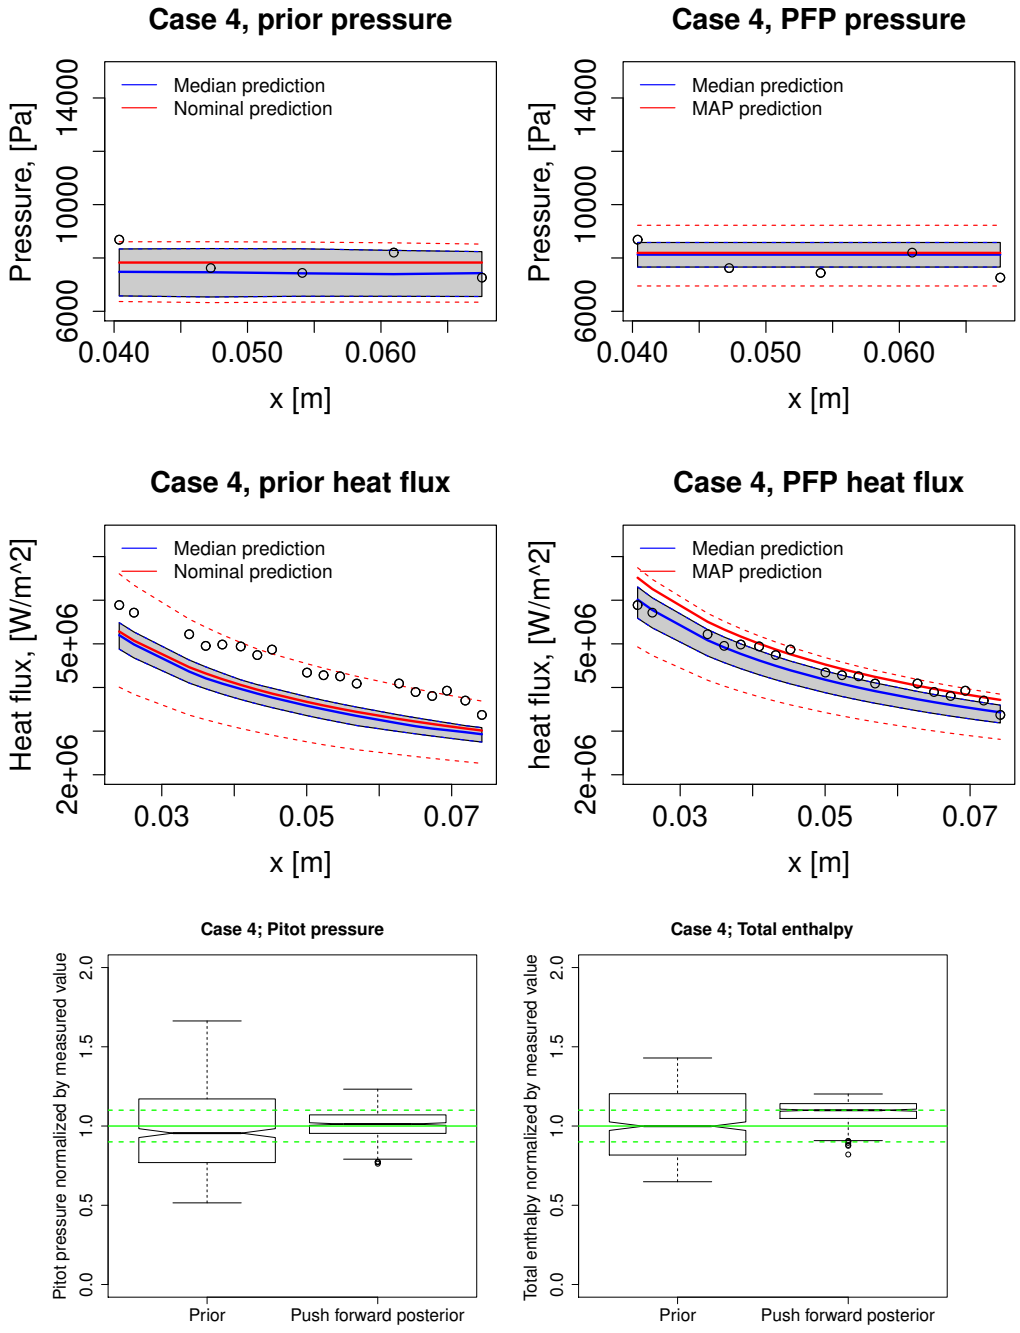

Figure 5: Top left: Predictions of pressure on the fore-cone using Θ samples drawn from the prior ($\pm 15\%$ bounds) for Case 4. The grey region is the prediction IQR and the red dotted lines the 5^{th}-and 95^{th}-percentile predictions. The blue line is the median prediction and solid red line the prediction using Θ_{MAP}. Top right: Pressure predictions from the PFP. Middle row: Same as the top row, but for heat flux. Bottom left: Box-and-whisker plot of predictions of Pitot pressure using Θ sampled from the prior and PFP and normalized by the measured value. Bottom right: Similar plots for h_0. The solid green horizontal line denotes a perfect match between predictions and observations. The dotted green horizontal lines are the $\pm 10\%$ limits.

4 The Quality of Case 4's Inference

Here we check the quality of the posterior PDF for Case 4 (Fig. 8 in the paper). In Fig. 5, we plot predictions from the push-forward prior and PFP simulations for pressure and heat flux on the fore-cone, as well as h_0 and P_{Pitot}. These predictions were performed using the statistical emulators seeded with 100 samples of Θ extracted from the posterior PDFs (plotted in Fig. 8 in the paper) and are limited to the fore-cone. The prior distribution is again uniform and spans $\pm 15\%$ around the nominal conditions in Table 1. We see that, post-calibration, the spread in the predictions has narrowed and the IQR contains about half the predictions. However, the calibration is still biased (the heat-flux measurements are greater than the median prediction in the blue line), though the bias is small for $p(x)$, and P_{Pitot}. Post-calibration, h_0 predictions are too large. Thus, the posterior PDF should result in a better match with the measurements in the attached, laminar fore-cone flow region, when simulated using SPARC. The more important question is whether it improves flow predictions over the entire double-cone.

References

[1] Nompelis, I., Candler, G. V., and Holden, M. S., "Effect of vibrational nonequilibrium on hypersonic double-cone experiments," *AIAA Journal*, Vol. 41, No. 11, 2003, pp. 2162–2169,
doi:10.2514/2.6834.

[2] Harvey, J., Holden, M., and Candler, G., *Validation of DSMC/Navier-Stokes Computations for Laminar Shock Wave/Boundary Layer Interactions Part 3*, 2003,
doi:10.2514/6.2003-3643.