ARTICLE

Single-cell RNA-seq reveals dynamic paracrine control of cellular variation

Alex K. Shalek[1,2,3]*, Rahul Satija[3]*, Joe Shuga[4]*, John J. Trombetta[3], Dave Gennert[3], Diana Lu[3], Peilin Chen[4], Rona S. Gertner[1,2], Jellert T. Gaublomme[1,2], Nir Yosef[3], Schraga Schwartz[3], Brian Fowler[4], Suzanne Weaver[4], Jing Wang[4], Xiaohui Wang[4], Ruihua Ding[1,2], Raktima Raychowdhury[3], Nir Friedman[5], Nir Hacohen[3,6], Hongkun Park[1,2,3], Andrew P. May[4] & Aviv Regev[3,7]

High-throughput single-cell transcriptomics offers an unbiased approach for understanding the extent, basis and function of gene expression variation between seemingly identical cells. Here we sequence single-cell RNA-seq libraries prepared from over 1,700 primary mouse bone-marrow-derived dendritic cells spanning several experimental conditions. We find substantial variation between identically stimulated dendritic cells, in both the fraction of cells detectably expressing a given messenger RNA and the transcript's level within expressing cells. Distinct gene modules are characterized by different temporal heterogeneity profiles. In particular, a 'core' module of antiviral genes is expressed very early by a few 'precocious' cells in response to uniform stimulation with a pathogenic component, but is later activated in all cells. By stimulating cells individually in sealed microfluidic chambers, analysing dendritic cells from knockout mice, and modulating secretion and extracellular signalling, we show that this response is coordinated by interferon-mediated paracrine signalling from these precocious cells. Notably, preventing cell-to-cell communication also substantially reduces variability between cells in the expression of an early-induced 'peaked' inflammatory module, suggesting that paracrine signalling additionally represses part of the inflammatory program. Our study highlights the importance of cell-to-cell communication in controlling cellular heterogeneity and reveals general strategies that multicellular populations can use to establish complex dynamic responses.

Variation in component molecules between individual cells[1–7] may have an important role in diversifying population-level responses[8–11], but also poses therapeutic challenges[4,5]. Although pioneering studies have explored heterogeneity within cell populations by focusing on small sets of preselected markers[1,2,4–6,8,12], single-cell genomics promises unbiased exploration of the molecular underpinnings and consequences of cellular variability[13–17].

We previously[16] used single-cell RNA-seq to identify substantial differences in messenger RNA (mRNA) transcript structure and abundance across 18 bone-marrow-derived mouse dendritic cells 4 h after stimulation with lipopolysaccharide (LPS, a component of Gram-negative bacteria). Many highly expressed immune response genes were distributed bimodally amongst single cells, originating, in part, from closely related maturity states and the variable activation of a key antiviral circuit. As these observations focused on a single pathogenic stimulus and time point, they raised several questions about the causes and roles of single-cell variability during the innate immune response. Examining the dynamics of cellular heterogeneity, its pathogen-specificity, and its intra- and intercellular control required new approaches to profile large numbers of cells from diverse conditions and genetic perturbations.

Here we use a microfluidic device to help prepare over 1,700 SMART-seq[15] single-cell RNA-seq libraries along time courses of bone-marrow-derived dendritic cells responding to different stimuli (Fig. 1 and Extended Data Fig. 1a). Combining computational analyses with diverse perturbations—including stimulation of individual cells in isolated, sealed microfluidic chambers and genetic and chemical alterations of paracrine signalling—we show how both antiviral and inflammatory response

modules in dendritic cells are controlled by positive and negative intercellular paracrine signalling that both promote and restrain variation.
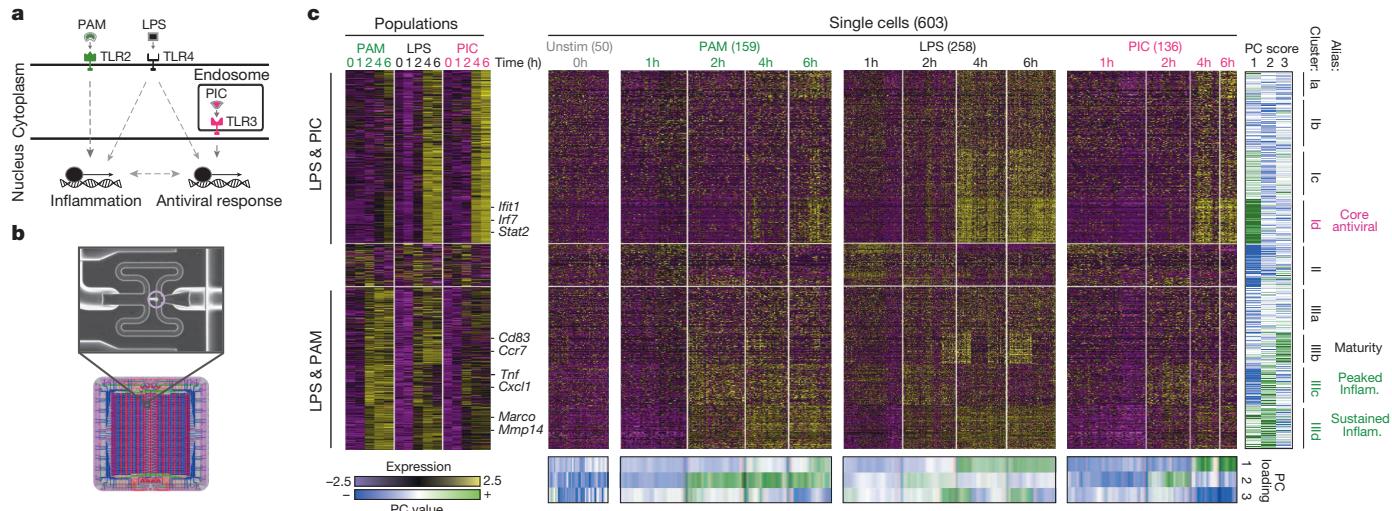
Microfluidics-based single-cell RNA-seq

We used the C_1 single-cell Auto Prep System (Fluidigm; Fig. 1b) and a transposase-based library preparation strategy to perform SMART-seq[15] (Supplementary Information) on 1,775 single dendritic cells, including both stimulation time courses (0, 1, 2, 4 and 6 h) for three pathogenic components[18] (LPS, PIC (viral-like double-stranded RNA), and PAM (synthetic mimic of bacterial lipopeptides)) and additional perturbations (Fig. 1, Extended Data Fig. 1 and Supplementary Information). For most conditions, we captured up to 96 cells (87 ± 8 (average ± s.d.)), and generated a matching population control (Fig. 1c, Supplementary Information and Supplementary Table 1). We prepared technically matched culture and stimulation replicates for the 2 h and 4 h LPS stimuli, and independent biological replicates for the unstimulated (0 h) and 4 h LPS experiments (Supplementary Information). We sequenced each sample to an average depth of 4.5 ± 3.0 million read pairs, as single-cell expression estimates stabilized at low read-depths[13,19] (Extended Data Fig. 2). The quality of our libraries was comparable to published SMART-seq data[15,16] (Extended Data Fig. 1b, Supplementary Tables 1 and 2). Overall, we successfully profiled 831 cells in our initial time courses and 944 cells in subsequent experiments (Extended Data Fig. 1a and Supplementary Tables 1 and 2). We excluded another 1,010 libraries with stringent quality criteria (Supplementary Information and Extended Data Fig. 1c).

Aggregated in silico, single-cell expression measurements agreed with the matching population controls ($R = 0.87 ± 0.05$), with correlations

Figure 1 | Microfluidic-enabled single-cell RNA-seq of dendritic cells stimulated with pathogenic components. a, Schematic of Toll-like receptor (TLR) sensing of PAM by TLR2, LPS by TLR4, and PIC by TLR3 (Supplementary Information). **b**, Microfluidic capture of a single dendritic cell (top, cell circled in purple) on a C_1 chip (CAD drawing, bottom). **c**, Time-course expression profiles for induced genes (rows) in dendritic cells (columns) at 0, 1, 2, 4 and 6 h post-stimulation with PAM (green), LPS (black), or PIC (magenta) within populations (left) and individual cells (right). Far right: gene projection scores onto the first three principal components (PCs) (columns); bottom: contributions of each cell (columns) to the first three PCs (rows).

plateauing once we had sampled ~30 cells (Supplementary Information, Extended Data Fig. 1d–g). Technical and biological replicates were reproducible (technical: aggregate $R > 0.90$, biological: aggregate $R > 0.87$; Extended Data Fig. 3) and our results were robust to variations in several aspects of sample preparation (Supplementary Information and Extended Data Fig. 1h–j). We removed 537 'cluster-disrupted' dendritic cells[16], a distinct subpopulation that matures as an artefact of isolation and culturing (Supplementary Information and Extended Data Fig. 4), retaining 1,238 dendritic cells for further analyses (Supplementary Tables 1 and 2).

Variability during immune responses

Principal components analysis (PCA) of gene expression profiles from all three time courses together showed that dendritic cells spread along a continuum of expression variation in each principal component (PC) (Fig. 1c and Extended Data Fig. 1k–n). For example, although PC1 distinguished early from late time points for each stimulus, its scores also varied substantially between cells within any single stimulus and time point (Fig. 1c and Extended Data Fig. 1k–n), suggesting that some cells were ahead of others, especially early (1–2 h).

Consistent with previous studies[18], pathogen-responsive genes partitioned into co-regulated modules based on their population-level expression profiles (Fig. 1c, left; Supplementary Information). Genes induced in cells stimulated with LPS or PIC (cluster I, Fig. 1c) were enriched for antiviral defence factors, including interferons and their targets (Bonferroni-corrected $P < 10^{-5}$), whereas genes induced in cells stimulated with LPS or PAM (cluster III, Fig. 1c) were enriched for inflammatory genes and NF-κB targets (Bonferroni-corrected $P < 10^{-6}$; Supplementary Table 3).

We used the single-cell gene expression profiles to partition these main clusters into finer modules (Fig. 1c, black lines, right; Supplementary Table 3; Supplementary Information) and applied a resampling method[20] (Supplementary Information, Extended Data Fig. 5d) to identify four modules significantly associated with the three major PCs (Fig. 1c): Cluster I_d (core antiviral module; enriched for annotated antiviral and interferon response genes; for example, *Ifit1*, *Irf7*; Bonferroni-corrected $P < 10^{-8}$, Supplementary Table 3, Fig. 1c and Extended Data Fig. 5a) was significantly associated with PC1; cluster III_c (peaked inflammatory module; showing rapid, yet transient, induction under LPS; for example, *Tnf*, *Il1a*, *Cxcl2*) and cluster III_d (sustained inflammatory module; exhibiting continued rise in expression under LPS; for example, *Mmp14*,

Marco, *Il6*) were associated with PC2; and cluster III_b ('maturity' module; containing markers of dendritic cell maturation; for example, *Cd83*, *Ccr7*, and *Ccl22*; Supplementary Information) was associated with PC3.

Digital and analogue expression variability

Genes from these four modules displayed distinct patterns of variation that changed with time and stimulus (Fig. 2a, Extended Data Figs 5 and 6). For example, early after LPS stimulation, core antiviral response genes were detectably expressed only in some cells (that is, were bimodal) (Fig. 2a, Extended Data Figs 5a and 6), but were turned on in most cells between 2 and 4 h (that is, became unimodal). In contrast, many peaked inflammatory genes were induced by LPS in all cells early, but were only detectable in some cells later (Fig. 2a, Extended Data Figs 5b and 6). Finally, sustained inflammatory genes were induced early in most cells and persisted at equal or elevated levels later (Fig. 2a, Extended Figs 5c and 6). Some variation patterns changed between stimuli (for example, peaked inflammatory genes remained detectably expressed in most cells late (6 h) in PAM), whereas other patterns were similar for distinct pathogens (for example, the antiviral modules I_a–I_d under LPS and PIC) (Figs 1 and 2a and Extended Data Fig. 5a–c).

As noted previously from single-cell quantitative real-time polymerase chain reaction (qRT–PCR) data[21], we distinguished two types of heterogeneity: (1) digital (on/off) variation, reflecting the percentage of cells detectably expressing a transcript; and (2) analogue variation, representing expression level variability among detectably expressing cells. Using the variance calculated over all cells as a metric of heterogeneity[6,16] conflates these two types of variation. We therefore explicitly modelled our data using three parameters (Fig. 2b and Extended Data Fig. 7): the mean (μ) and variance ($σ^2$) of a gene's expression among detectably expressing cells, and the fraction of detectably expressing cells (α)[21]: in this scheme, $σ^2$ and α signify analogue and digital variation, respectively.

We computed α based on a fixed threshold for appreciable expression ($\ln(\text{TPM} + 1) > 1$, Supplementary Information and Extended Data Fig. 7a, f), and then estimated μ and $σ^2$ across appreciably expressing cells. This three-parameter model effectively described most (91%) of our single-cell data (Fig. 2c, d, Supplementary Information and Extended Data Fig. 7b). Our data did not support fitting with either a single log-normal or a mixture of two, fully parameterized lognormals (modelling high and low expression states; Supplementary Information and Extended Data Fig. 7c–e). Computed α values were consistent between technical

Figure 2 | Time-dependent behaviours of single cells. a, Single-cell expression distributions for three genes at each time point after stimulation with PAM (top, green), LPS (middle, black), or PIC (bottom, magenta). Distributions are scaled to have the same maximum height. Individual cells are plotted as bars underneath each distribution. **b**, Three parameters describing single-cell gene expression distributions: μ (green) and σ^2 (gold), the mean and variance of RNA abundance in detectably expressing cells, respectively, and α (blue), the fraction cells with detectable expression (at $\ln(\text{TPM} + 1) > 1$). **c**, Examples of fit (grey) and measured *Tnf* expression distributions (black). **d**, The values of μ, σ^2, and α (y axes, left to right) computed for *Tnf* at each time point (x axis). Units for μ and σ^2 are $\ln(\text{TPM} + 1)$. **e**, Maximum likelihood estimate α (α_{MLE}). Shown are the likelihood functions (dotted blue line) for *Tnf* (matching **c**) used to determine α_{MLE} (red line; vertical black line: nominal α; Supplementary Information).

and biological replicates, but μ and σ^2 estimates were reproducible only when genes were expressed in at least 10 or 30 cells, respectively (Supplementary Note, Supplementary Information, Extended Data Figs 2c–e, 7g and 8).

Our nominal α estimates are likely deflated due to the detection limits of single-cell RNA-seq. Indeed, we observe higher α values when examining our existing RNA fluorescence *in situ* hybridization (RNA-FISH) data[16] (Extended Data Fig. 6g–j). By comparing our single-cell RNA-seq and RNA-FISH, we estimate that the transcript detection efficiency for our single-cell RNA-seq is ~20%, consistent with previous reports[14,22]. We and others[15,23] have also observed a strong relationship between the average expression of a gene and its probability of detection (Extended Data Fig. 7h). We thus used a conservative null model, where this relationship results solely from technical limitations (Supplementary Information, Extended Data Fig. 7h), and determined the maximum likelihood estimate of α (α_{MLE}) for each gene after correcting for this relationship (Fig. 2e, Extended Data Fig. 7j–l and Supplementary Information). From this analysis, we estimate that ~45% of core antiviral genes and 30% of peaked inflammatory genes are significantly bimodal in at least one measured time point in the LPS response (likelihood ratio test (LRT), Bonferroni-corrected $P < 0.01$; Supplementary Information and Extended Data Fig. 7i).

Chromatin mark levels correlate with α
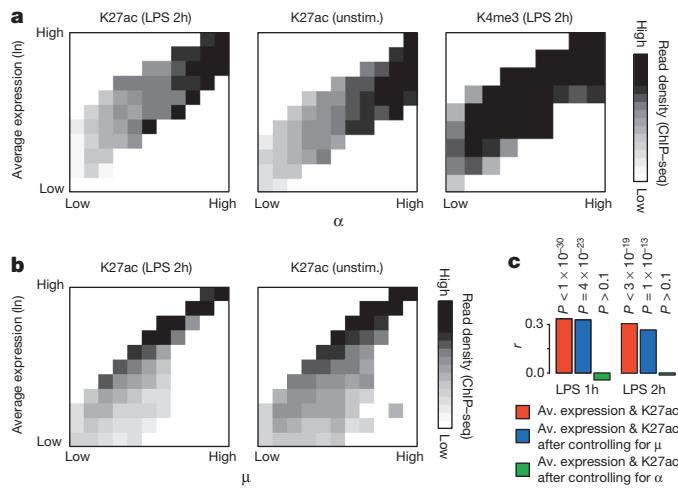
As the presence of a chromatin mark is, by definition, discrete in a single cell, we reasoned that population ChIP-seq profiles of active histone marks (for example, histone 3 lysine 27 acetylation (H3K27ac)) should more closely reflect the fraction of cells with detectable transcripts (α) than population-level expression. Supporting this hypothesis, the observed α for a gene was strongly correlated (mean R for binned data = 0.89; Supplementary Information) to its promoter-associated ChIP-seq

density (collected under identical conditions[24]), even within a fixed population expression range (Fig. 3a top/middle, rows). In contrast, a gene's population-level expression was not correlated (mean R for binned data = −0.02) to H3K27ac promoter levels within a fixed α range (Fig. 3a top, middle; columns). We note that H3K27ac and population-level expression remained correlated within a fixed range of μ (instead of α, Fig. 3b). A partial correlation analysis focussed on either all immune response genes or 'bimodal' genes (LRT, $P < 0.01$) yielded similar results ($P > 0.1$, after controlling for α, Fig. 3c). Digital variation did not correlate with histone 3 lysine 4 trimethylation (H3K4me3) levels (Fig. 3a, bottom), in line with previous observations[24] that H3K4me3 is not as tightly correlated with active transcription. Emerging single-cell epigenomic technologies[25] should help to further explore this relationship.

Dynamic responses via shifts in α and μ

An average (population) increase in the expression of bimodally expressed genes may represent changes in the amount of transcript made by expressing cells (shifts in μ), the proportion of expressing cells (shifts in α), or both. For each pair of consecutive time points, we examined the proportion of genes in each module with a significant change in: (1) μ (Wilcoxon rank-sum test); (2) α (LRT, controlling for the aforementioned confounding relationship between average expression and detection efficiency, Supplementary Information); or (3) both. Given our limitations in estimating α and μ, we only considered genes that were annotated as bimodal in at least one time point in the relevant time course and expressed in at least 10 cells in both time points (Supplementary Information). We excluded the unstimulated time point when most immune response genes were not yet expressed.

Under LPS stimulation, core antiviral and sustained inflammatory genes had the strongest increases in α (alone or with μ) at early time points (Fig. 3d, top; Extended Data Fig. 5e, f), and transitioned to high
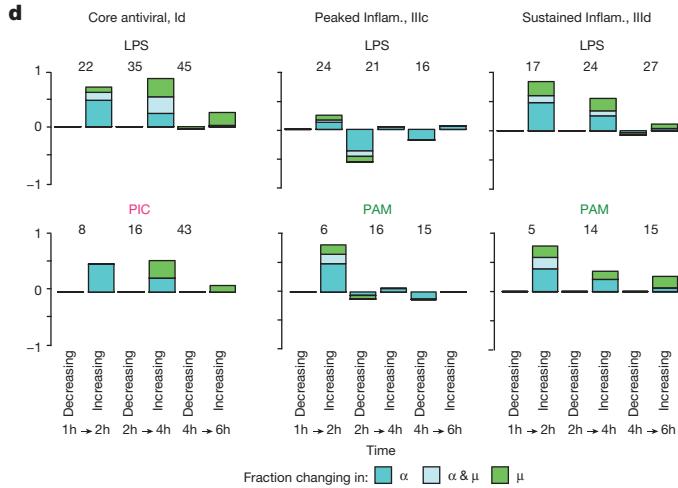
Figure 3 | Dynamic changes in variation during stimulation. a, b, The relationship between expression and H3K27ac binding depends on α (**a**), but not on μ (**b**). Plots show average promoter read density (black high; white low; scale bar, bottom) for H3K27ac in LPS 2 h (**a, b**, left) and unstimulated cells (**a**, middle; **b**, right), or H3K4me3 in 2 h LPS (**a**, right) in genes corresponding to each of 10 quantile bins of population expression (*y* axis) and each of 10 quantile bins of α (**a**, *x* axis) or μ (**b**, *x* axis) (Supplementary Information). **c**, Bar plots showing *P* values of correlation between average expression levels and K27ac only for immune response genes either as is (red) or when controlling for μ (blue) or α (green). Matching R values for all genes: 0.29 (LPS 1 h, as is),

0.18 (LPS 1 h, controlling for μ), 0.06 (LPS 1 h, controlling for α), 0.33 (LPS 2 h, as is), 0.23 (LPS 2 h, controlling for μ), 0.08 (LPS 2 h, controlling for α) **d**, Dynamic changes in α and μ in each module. Bar plots showing, for each module in select conditions (annotated on top), the fraction of genes (*y* axis) with a significant change only in α (*P* < 0.01, likelihood ratio test, blue), only in μ (*P* < 0.01, Wilcoxon test, green), or in both (each test independently, light blue), at each transition (*x* axis). The number of genes over which the proportion is calculated is marked on top of each pair of bars (Supplementary Information, Extended Data Fig. 5f).

and unimodal expression by 4 h (Figs 1 and 2). In contrast, α decreased at later time points for peaked inflammatory genes, especially from 2 to 4 h (Fig. 3d, middle; Extended Data Fig. 5f). The temporal patterns in core antiviral gene activation were shared between LPS and PIC. However, unlike in LPS, peaked inflammatory gene expression did not diminish under PAM, and we did not observe statistically significant decreases in α at later time points (Fig. 3d). These coherent shifts suggest that variability reflects regulated immune response phenomena, rather than unconstrained stochastic transcription.

Intercellular determinants of variation

Both differences in intracellular components[1–4] and changes in the cellular microenvironment[7,26] can affect heterogeneity. In particular, slow diffusion of cytokines and chemokines could lead to local variation in intercellular signals. As the core antiviral module is enriched for targets of IFN-β, we speculated that upstream variability in IFN-β exposure may drive its heterogeneity (median α = 0.52; 30% of genes significantly bimodal, *P* < 0.01, LRT, Extended Data Fig. 9), and thus

profiled cells 2 h after IFN-β stimulation. Supporting our hypothesis, cells stimulated with IFN-β for 2 h exhibited sharply reduced digital variation in the core antiviral module (Fig. 4a, median α = 0.82; 7% of genes significantly bimodal).

Precocious expressers of antiviral genes

We next explored the cellular source of interferon in the native LPS response. At 2 h following LPS, *Ifnb1* was bimodally expressed (*P* < 10⁻⁴, LRT) and correlated with the expression of the core antiviral module (Extended Data Fig. 9a, d, e). This observation, together with the suppression of digital variation under an IFN-β stimulus, suggested that, in response to LPS, a few cells may first produce (Extended Data Fig. 9d) and secrete a wave of interferon, leading to a gradual coordination of the core antiviral module at later time points via paracrine signalling.

To test this hypothesis, we computed a core antiviral activation score (Supplementary Information) for each cell and compared scores across the LPS time course (Fig. 4b, Extended Data Figs 9e, f and 10a and Supplementary Information). Although most cells activated the module

Figure 4 | IFN-β feedback drives heterogeneity in the expression of core antiviral targets. a, Single-cell expression distributions for *Rsad* (top) and *Stat2* (bottom) after stimulating with LPS (left, black) or IFN-β (right, magenta) for 2 h. **b**, The core antiviral score (*y* axis; Supplementary Information, Extended Data Figs 9f and 10a) for each LPS-stimulated cell at each time point

(*x* axis) and cells simulated for 2 h with IFN-β (on far right). Two precocious cells (yellow asterisk) have unusually high antiviral scores at 1 h LPS. **c**, RNA-FISH confirms the presence of rare precocious responders (arrow; yellow asterisk), positive for both *Ifnb1* (magenta) and *Ifit1* (cyan) 1 h after LPS stimulation. Grey, cell outlines. Scale bar represents 25 μm.

between 2 h and 4 h, we discovered two cells with strong core antiviral activation at 1 h (Fig. 4b, c, Extended Data Fig. 9f, i, yellow asterisks). We verified the existence and scarcity of these precocious cells 1 h after LPS stimulation by RNA-FISH (Fig. 4c, Supplementary Information); here, appreciable *Ifit1* and *Ifnb1* co-expression was detected in only 0.8% of cells (23 of 2,960, mRNA count ≥ 5 copies, $P = 2 \times 10^{-28}$, proportion test). These precocious cells were indistinguishable from the others except in their expression of the ~100 core antiviral genes (Extended Data Fig. 9j, k). We observed similar early responding cells following PIC or PAM stimulation (Extended Data Fig. 9f, h and 10a).

Although these precocious cells are reminiscent of the 'sentinels' that have been reported in viral infections and stimulations of fibroblasts[27,28] (Supplementary Note, Supplementary Information), we note that, in those studies, variable response may be partially due to differences in cells' ability to sense and respond to the primary stimulus (for example, due to lack of viral sensing or replication). In contrast, all dendritic cells rapidly sense and respond to LPS, as evidenced by the unimodal activation of peaked inflammatory genes at early time points (Extended Data Figs 5b–d, 10a; Fig. 2a, *Tnf*).
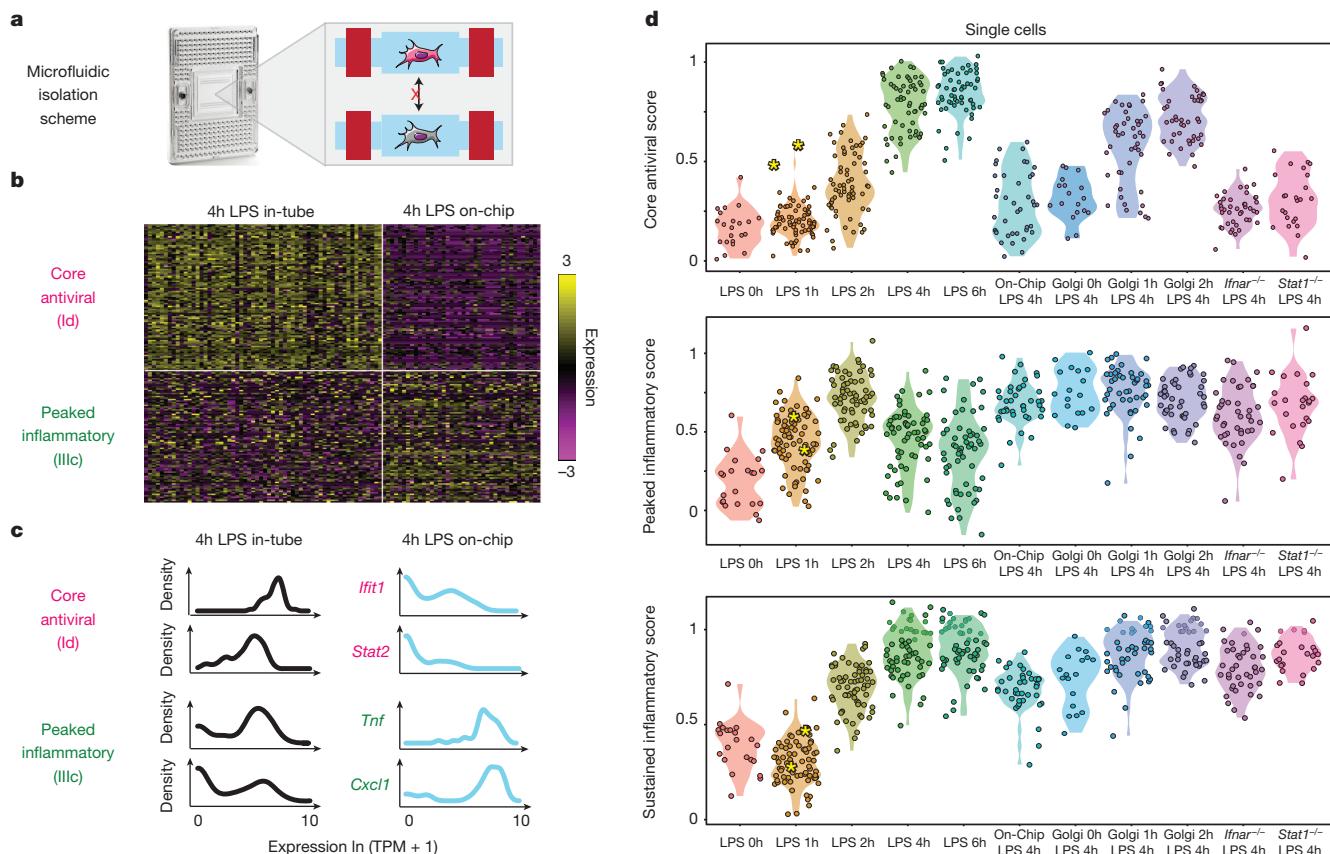
Intercellular communication and variation

To examine whether the rare precocious cells were required for coordinating the core antiviral response, we developed an approach to stimulate cells in the absence of cell-to-cell communication. Modifying the standard C_1 workflow, we captured individual unstimulated dendritic cells in a C_1 chip (Supplementary Information), washed in LPS-containing

media, and then immediately sealed each microfluidic chamber to isolate stimulated cells individually for 4 h (on-chip stimulation, Supplementary Information, Fig. 5a). Key experimental conditions, including cell density, were similar between the in-tube and on-chip experiments (Supplementary Information).

In the absence of cell-to-cell communication, core antiviral module genes were bimodally expressed (Fig. 5), with only 8 cells (20%) exhibiting weak activation of the core antiviral module at 4 h (Fig. 5b–d, Extended Data Fig. 9e), probably mimicking the precocious cells observed in-tube at 1 h. This observation suggests an approximate bound for the number of cells capable of autonomously inducing a response by 4 h. Removing cell-to-cell communication also downregulated the expression of maturation markers in all cells and some of the sustained inflammatory genes (Extended Data Fig. 10a), although other key inflammatory genes were unaffected.

Surprisingly, blocking intercellular communication also sharply altered the single-cell expression of peaked inflammatory genes (Fig. 5b–d). Genes encoding key inflammatory cytokines (for example, *Tnf*, *Cxcl1*) switched from bimodal ($\alpha = 0.77, 0.56$, respectively) to unimodal ($\alpha = 1.0$, 0.91; LRT for corresponding α_{MLE}: $P < 10^{-4}$, $P < 10^{-13}$, respectively) expression on-chip (Fig. 5b, c). Indeed, a large portion of the peaked inflammatory genes that were bimodal (LRT $P < 0.01$) after a 4 h LPS stimulation in-tube shifted to unimodal expression on-chip (Extended Data Fig. 10a, b; $P < 0.01$, hypergeometric test), indicating that cell-to-cell signalling is required for dampening the peaked inflammatory program at later time points following LPS. The opposite behaviours

Figure 5 | Microfluidic blocking of cell-to-cell signalling affects response heterogeneity in the core antiviral and peaked inflammatory modules.
a, Experimental blocking of cell-to-cell communication. Left: C_1 chip; right: actuation of microfluidic valves (red bars), following on-chip LPS stimulation, isolates individual cells in sealed chambers, preventing intercellular signalling. **b**, Expression of the genes (rows) in the core antiviral (I_d, top rows) and peaked inflammatory (III_c, bottom rows) modules in single cells (columns) from the in-tube (left) and on-chip (right) stimulations. **c**, Gene expression

distributions for representative genes from the core antiviral (top) and peaked inflammatory (bottom) modules in the in-tube (left, black) or on-chip (right, blue) 4 h LPS stimulation. **d**, Violin plots of, top to bottom, the core antiviral (Supplementary Information, top), peaked inflammatory (middle), and sustained inflammatory (bottom) scores for individual cells from the stimulation conditions listed on the x axis. Yellow asterisks: the two precocious cells from Fig. 4 (Extended Data Fig. 10a).

of the core antiviral and peaked inflammatory modules indicate that intercellular communication can have opposing effects on variation for different gene modules within the same cell.

IFN-β and peaked inflammatory genes

On-chip isolation conflates the effects of different paracrine signals and the loss of cell-to-cell contact. To distinguish these situations, we first profiled dendritic cells from Interferon receptor knockout mice ($Ifnar1^{-/-}$). As expected, and consistent with previous findings[16], antiviral gene expression was undetectable at 4 h in all $Ifnar1^{-/-}$ dendritic cells, implying that even the precocious cells require autocrine interferon feedback to activate and sustain their core antiviral responses (Extended Data Fig. 10g). This is further supported by the decoupling of the expression of $Ifnb1$ and the core antiviral module in $Ifnar1^{-/-}$ dendritic cells stimulated with LPS for 2 h (Extended Data Fig. 9e).

Removal of interferon signalling also strongly affected the peaked inflammatory module: after 4 h of LPS stimulation, $Ifnar1^{-/-}$ cells showed a similar increase in the fraction of activated cells as was seen in the on-chip experiment (Fig. 5d, Extended Data Fig. 10a, d, g), suggesting that the absence of interferon signalling, rather than changes in cell-to-cell contact[29], was the major driver. Furthermore, dendritic cells lacking $Stat1$, a gene encoding a key transcription factor mediating interferon responses[24], also exhibited increased activation and decreased digital variation in peaked inflammatory genes ($P < 0.01$; hypergeometric test; Fig. 5d and Extended Data Fig. 10a, e, g, i). Conversely, the sustained inflammatory module was not appreciably affected by the absence of interferon signalling (Fig. 5d and Extended Data Fig. 10a, g), implying a different mechanism for its downregulation on-chip.

Second paracrine wave for downregulation

Interferon response targets can cross-inhibit inflammatory gene expression either through the direct formation of repressive complexes, for example, the STAT1-inclusive ISGF-3, or by inducing the production of anti-inflammatory cytokines[30]. The few cells with on-chip antiviral activation exhibited no change in peaked inflammatory gene expression (Fig. 5b). This suggests that the repression of peaked inflammatory genes, unlike antiviral activation, is not directly downstream of IFN-β signalling, but rather may be mediated by a second IFN-β/STAT1-dependent paracrine signal. Peaked induction through two asynchronous paracrine signals is reminiscent of the activation and contraction of keratinocytes following wounding and immune infiltration, respectively[31].

To test this hypothesis further, we added brefeldin A (GolgiPlug), a secretion inhibitor, either simultaneously with LPS (0 h) or at 1 or 2 h after stimulation, and measured single-cell RNA-seq profiles at 4 h (Fig. 5d, Extended Data Fig. 10a–c). Inhibiting secretion at the time of LPS addition strikingly dampened the antiviral response, similar to the on-chip experiment. However, adding brefeldin A at 2 h did not affect the activation of the core antiviral module and adding it at 1 h had only a modest effect. This indicates that the first hour represents the crucial paracrine window for this response. In contrast, for the peaked inflammatory module, addition at each of the three time points resulted in the module remaining aberrantly activated at 4 h, as on-chip. Collectively, these experiments show that paracrine interferon signalling events before the 1 h time point are crucial for antiviral activation, whereas subsequent, separate signalling is responsible for the desynchronized dampening of peaked inflammatory gene expression (Supplementary Note, Supplementary Information).

Discussion

Here we have analysed how gene expression variation between individual dendritic cells changes with stimulus and time to dissect the regulation of heterogeneity across this immune response. Our statistical analysis reveals that changes in digital (on/off) variation can encode a diversity of temporal response profiles (Fig. 3d, Extended Data Fig. 5f). For example, late-induced core antiviral genes are very weakly expressed early, on average, but are highly expressed in a few precocious cells; the progressive dampening of peaked inflammatory genes originates, in part, from changes in the fraction of cells detectably expressing these transcripts, rather than a uniform, gradual decrease in their expression in all cells.

Such complex average responses can be generated not only through intricate intracellular circuits in each cell, but also through intercellular communication between cells, as we show for both modules. For example, we uncovered a small number of precocious cells that express $Ifnb1$ and core antiviral genes as early as 1 h after LPS stimulation, and through the secretion of IFN-β, help activate core antiviral genes in other cells to coordinate the population response. These cells are indistinguishable from the rest, except in their expression of the core antiviral module (Extended Data Fig. 9j, k), and yet are crucial for an efficient and timely population response (Supplementary Note, Supplementary Information).

IFN-β signalling also dampens a subset of induced inflammatory genes at later time points, and our brefeldin A (GolgiPlug) experiments suggest that a secondary, IFN-β-dependent signal, is involved (Extended Data Fig. 10j, k). This is consistent with a model in which IFN-β secreted by a few cells induces the expression and secretion of secondary anti-inflammatory cytokines from a subset of cells, which, in turn, attenuate the peaked inflammatory responses of their neighbours. Computational analyses, genetic perturbations and recombinant cytokine experiments suggest that IL-10 may be involved in this second wave of negative signalling (Extended Data Fig. 10h, Supplementary Table 4), but further experiments are needed to fully elucidate the mechanism (Supplementary Note, Supplementary Information). One involved component may be the RNA degradation factor ZFP36 (TTP), whose targets are enriched in the peaked inflammatory module[32].

The ability of precocious cells to influence others via paracrine signalling may be an efficient strategy for quorum sensing[33], but also may be perilous. If the activation threshold is too low, a few stochastically responding cells could induce an inappropriate immune response. Indeed, this is observed in autoimmune diseases like systemic lupus erythematosus (SLE), in which excess IFN-β production potentiates auto-reactive dendritic cell activation[34,35]. In contrast, excessively stringent thresholds may limit rapid responses to a viral infection, or the dampening of chronic inflammation (for example, in rheumatoid arthritis or ulcerative colitis[30,35]). Thus, individual cells probably place tight controls on the regulation of key cytokines, preferring different induction strategies under different stimuli to maximize the balance between responsiveness and control. Indeed, similar population-level $Ifnb1$ expression in LPS/PIC (Extended Data Fig. 9c) stems from different underlying phenomena: a substantial fraction of cells express the $Ifnb1$ transcript moderately at 2 h LPS ($\alpha = 0.35$, $\mu = 5.1$), whereas just a few cells express $Ifnb1$ very highly at 2 h PIC ($\alpha = 0.07$, $\mu = 6.31$; uncorrelated with the cell's activation of the antiviral response[26,27]: Extended Data Fig. 9e).

Using microfluidics, we achieved the statistical power needed to track transcriptome-wide changes in single-cell expression variation across a variety of conditions, as well as to identify functionally important, rare responses. Microfluidics also allowed us to finely control the stimulation of our cells. Similar and improved techniques will be essential for characterizing other rare sub-populations, such as cancer stem cells, and for studying heterogeneous clinical samples and tissues. Further innovations in massively parallel manipulation and profiling of single cells will continue to improve our understanding of the rich diversity in, and dynamic functional communities that constitute, multicellular populations.

METHODS SUMMARY

Bone-marrow-derived mouse dendritic cells were prepared as previously described[18] and stimulated with pathogenic stimuli for specified time periods. The C_1 Single-Cell Auto Prep System (Fluidigm) was used to perform SMARTer (Clontech) whole transcriptome amplification (WTA)[15,16,19] on up to 96 individual cells. WTA products were then converted to Illumina sequencing libraries using Nextera XT (Illumina)[15]. RNA-seq libraries were also made from 10,000 cells from each parent population (population control). Each sample was sequenced on an Illumina HiSeq 2000 or 2500, and expression estimates (transcripts per million; TPM) for all UCSC-annotated

mouse genes were calculated using RSEM[36]. Data were further analysed as described in the Supplementary Information. Additional experiments were performed using RNA-FISH (Panomics), on-chip isolated stimulation, knockout mice, secretion blockers (GolgiPlug, BD Biosciences), protein synthesis blockers (cycloheximide, Sigma), and recombinant cytokines. Full Methods and any associated references are provided in the Supplementary Information.

Online Content Any additional Methods, Extended Data display items and Source Data are available in the online version of the paper; references unique to these sections appear only in the online paper.
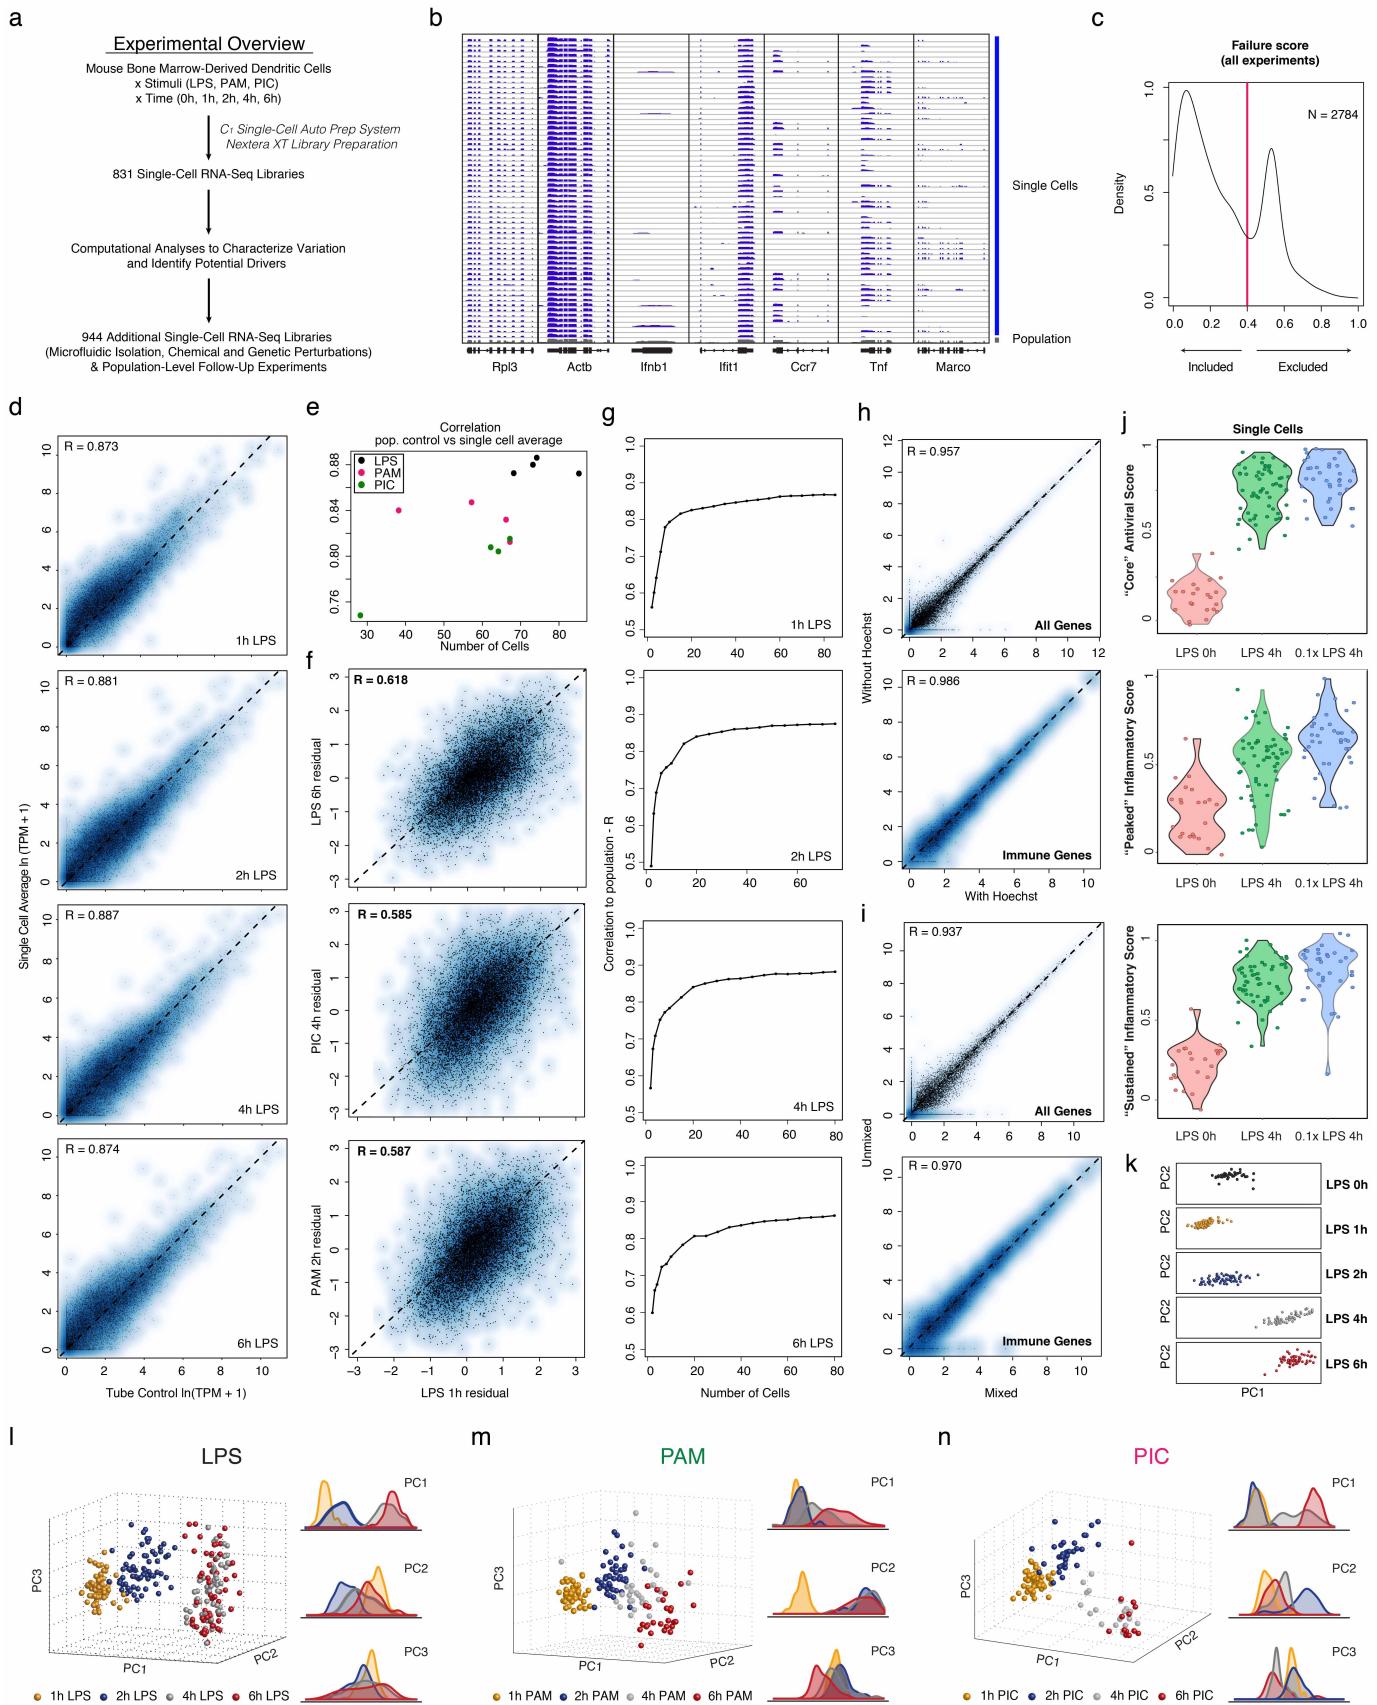
1. Tay, S. et al. Single-cell NF-κB dynamics reveal digital activation and analogue information processing. Nature 466, 267–271 (2010).
2. Raj, A. & Van Oudenaarden, A. Single-molecule approaches to stochastic gene expression. Ann. Rev. Biophys. 38, 255–270 (2009).
3. Slack, M. D., Martinez, E. D., Wu, L. F. & Altschuler, S. J. Characterizing heterogeneous cellular responses to perturbations. Proc. Natl Acad. Sci. USA 105, 19306–19311 (2008).
4. Sharma, S. V. et al. A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. Cell 141, 69–80 (2010).
5. Spencer, S. L., Gaudet, S., Albeck, J. G., Burke, J. M. & Sorger, P. K. Non-genetic origins of cell-to-cell variability in TRAIL-induced apoptosis. Nature 459, 428–432 (2009).
6. Taniguchi, Y. et al. Quantifying E. coli proteome and transcriptome with single-molecule sensitivity in single cells. Science 329, 533–538 (2010).
7. Loewer, A. & Lahav, G. We are all individuals: causes and consequences of non-genetic heterogeneity in mammalian cells. Curr. Opin. Genet. Dev. 21, 753–758 (2011).
8. Feinerman, O. et al. Single-cell quantification of IL-2 response by effector and regulatory T cells reveals critical plasticity in immune response. Mol. Syst. Biol. 6, 437–453 (2010).
9. Veening, J.-W., Smits, W. K. & Kuipers, O. P. Bistability, epigenetics, and bet-hedging in bacteria. Ann. Rev. Microbiol. 62, 193–210 (2008).
10. Fang, M., Xie, H., Dougan, S. K., Ploegh, H. & Van Oudenaarden, A. Stochastic cytokine expression induces mixed T helper cell states. PLoS Biol. 11, e1001618 (2013).
11. Chalancon, G. et al. Interplay between gene expression noise and regulatory network architecture. Trends Genet. 28, 221–232 (2012).
12. Bendall, S. C. et al. Single-cell mass cytometry of differential immune and drug responses across a human hematopoietic continuum. Science 332, 687–696 (2011).
13. Hashimshony, T., Wagner, F., Sher, N. & Yanai, I. CEL-Seq: single-cell RNA-Seq by multiplexed linear amplification. Cell Rep. 2, 666–673 (2012).
14. Islam, S. et al. Characterization of the single-cell transcriptional landscape by highly multiplex RNA-seq. Genome Res. 21, 1160–1167 (2011).
15. Ramsköld, D. et al. Full-length mRNA-Seq from single-cell levels of RNA and individual circulating tumor cells. Nature Biotechnol. 30, 777–782 (2012).
16. Shalek, A. K. et al. Single-cell transcriptomics reveals bimodality in expression and splicing in immune cells. Nature 498, 236–240 (2013).
17. Tang, F. et al. mRNA-Seq whole-transcriptome analysis of a single cell. Nature Methods 6, 377–382 (2009).
18. Amit, I. et al. Unbiased reconstruction of a mammalian transcriptional network mediating pathogen responses. Science 326, 257–263 (2009).
19. Wu, A. R. et al. Quantitative assessment of single-cell RNA-sequencing methods. Nature Methods 11, 41–46 (2014).
20. Storey, J. D. & Tibshirani, R. Statistical significance for genomewide studies. Proc. Natl Acad. Sci. USA 100, 9440–9445 (2003).
21. McDavid, A. et al. Data exploration, quality control and testing in single-cell qPCR-based gene expression experiments. Bioinformatics 29, 461–467 (2013).
22. Islam, S. et al. Highly multiplexed and strand-specific single-cell RNA 5′ end sequencing. Nature Protocols 7, 823–828 (2012).
23. Marinov, G. K. et al. From single-cell to cell-pool transcriptomes: stochasticity in gene expression and RNA splicing. Genome Res. 24, 496–510 (2014).
24. Garber, M. et al. A high-throughput chromatin immunoprecipitation approach reveals principles of dynamic gene regulation in mammals. Mol. Cell 47, 810–822 (2012).
25. Weibrecht, I. et al. Visualising individual sequence-specific protein-DNA interactions in situ. New Biotechnol. 29, 589–598 (2012).
26. Lee, T. K. et al. A noisy paracrine signal determines the cellular NFκB response to lipopolysaccharide. Sci. Signal. 2, ra65 (2009).
27. Rand, U. et al. Multi-layered stochasticity and paracrine signal propagation shape the type-I interferon response. Mol. Syst. Biol. 8, 584 (2012).
28. Zhao, M., Zhang, J., Phatnani, H., Scheu, S. & Maniatis, T. Stochastic expression of the interferon-β gene. PLoS Biol. 10, e1001249 (2012).
29. Snijder, B. et al. Population context determines cell-to-cell variability in endocytosis and virus infection. Nature 461, 520–523 (2009).
30. Benveniste, E. N., Qin, H. & Type, I. Interferons as anti-inflammatory mediators. Sci. STKE 416, pe70 (2007).
31. Freedberg, I. M., Tomic-Canic, M., Komine, M. & Blumenberg, M. Keratins and the keratinocyte activation cycle. J. Invest. Dermatol. 116, 633–640 (2001).
32. Rabani, M. et al. Metabolic labeling of RNA uncovers principles of RNA production and degradation dynamics in mammalian cells. Nature Biotechnol. 29, 436–442 (2011).
33. Waters, C. M. & Bassler, B. L. Quorum sensing: cell-to-cell communication in bacteria. Ann. Rev. Cell Dev. Biol. 21, 319–346 (2005).
34. Banchereau, J., Pascual, V. & Type, I. Interferon in systemic lupus erythematosus and other autoimmune diseases. Immunity 25, 383–392 (2006).
35. Hall, J. C. & Rosen, A. Type I interferons: crucial participants in disease amplification in autoimmunity. Nature Rev. Rheumatol. 6, 40–49 (2010).
36. Li, B. & Dewey, C. N. RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. BMC Bioinformatics 12, 323 (2011).

a Experimental Overview

Mouse Bone Marrow-Derived Dendritic Cells
x Stimuli (LPS, PAM, PIC)
x Time (0h, 1h, 2h, 4h, 6h)

C_1 Single-Cell Auto Prep System
Nextera XT Library Preparation

831 Single-Cell RNA-Seq Libraries

Computational Analyses to Characterize Variation
and Identify Potential Drivers

944 Additional Single-Cell RNA-Seq Libraries
(Microfluidic Isolation, Chemical and Genetic Perturbations)
& Population-Level Follow-Up Experiments

Extended Data Figure 1 | Single-cell RNA-seq of hundreds of dendritic cells.
a, Overview of experimental workflow. **b**, Shown are read densities for seven representative genes (two housekeeping genes (*Rpl3* and *Actb*) and five immune response genes (*Ifnb* , *Ifit1*, *Ccr7*, *Tnf*, *Marco*)) across 60 single cells (blue) and one population control of 10,000 cells (grey; bulk population) after a 4 h LPS treatment. **c**, Distribution of failure scores for all single cells. Single cells with failure scores above 0.4 were discarded (see Supplementary Information). **d–g**, Comparisons of expression estimates for the average single cell and the bulk population. **d**, Scatter plots showing for each gene the relation between the average single-cell expression (*y* axis) and bulk population level expression (*x* axis) for each of four time points following LPS stimulation (1, 2, 4 and 6 h, left to right). **e**, The Pearson correlation coefficients (*y* axis) of each comparison, as in **d**, for each of the time points and stimuli presented in Fig. 1, as a function of the number of cells captured in the respective experiment (*x* axis). **f**, Scatter plots showing the residual (population-level expression minus the single cell average) in a LPS 1 h experiment (*x* axis) versus the residual in each of 3 other experiments (*y* axis, left to right): LPS 6 h, PIC 4 h and PAM 2 h. **g**, The Pearson correlation coefficient (*y* axis) between the bulk population level expression and the single-cell expression average when a different number of sub-sampled cells (*x* axis) is included in the single-cell
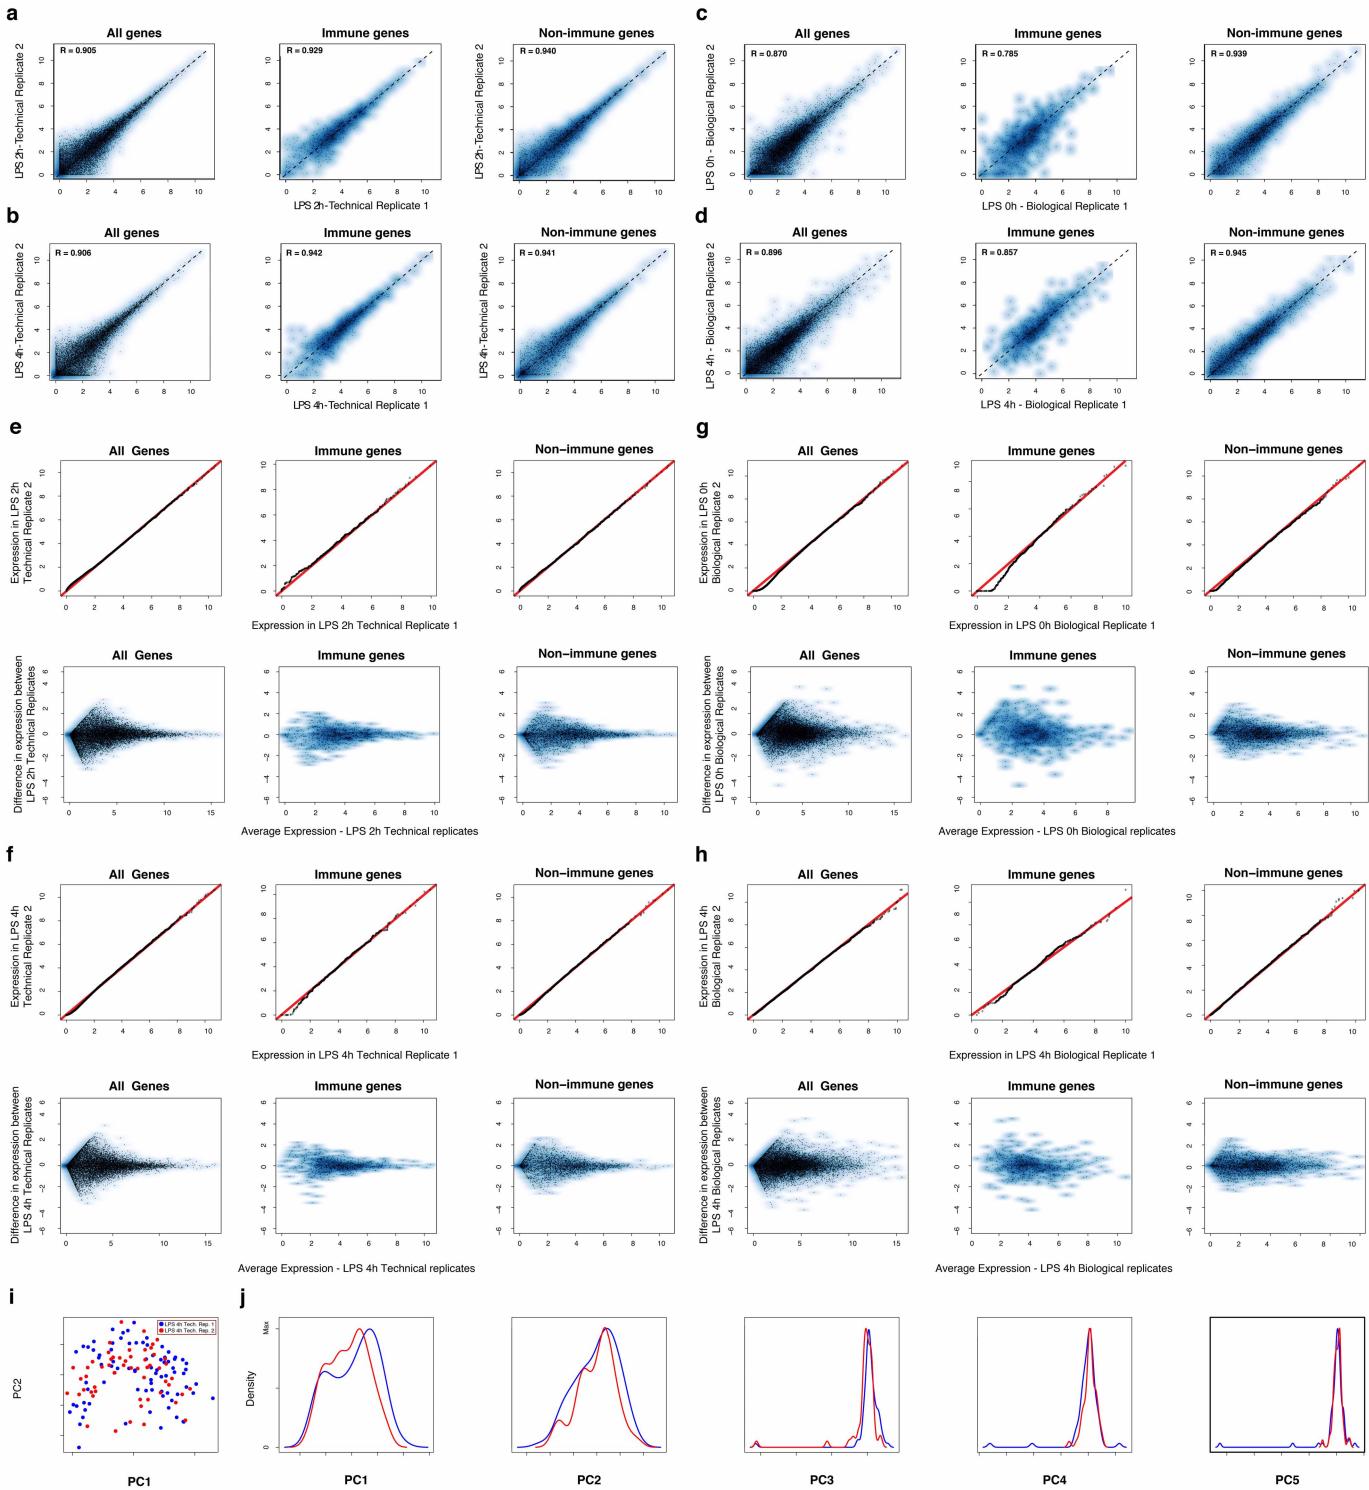
average. **h**, **i**, Effects of Hoechst dye and periodic mixing on mRNA expression. **h**, Comparable expression levels after 4 h LPS with the addition of small amounts of Hoechst to aid in cell counting (*x* axis) and when no dye is used (*y* axis), when looking at all genes (left) or only immune response genes (right). **i**, Comparable expression levels after 4 h LPS with hourly mixing (*x* axis) or with no mixing (*y* axis), when looking at all genes (left) or only immune response elements (right). **j**, Core antiviral, peaked inflammatory, and sustained inflammatory module activation scores for a 0.1× LPS stimulation. Shown are violin plots of the scores (*y* axis) for the core antiviral (Supplementary Information, top), peaked inflammatory (Supplementary Information, middle), and sustained inflammatory modules (Supplementary Information, bottom) for each cell in (left to right): LPS 0 h, 1× (100 ng ml^{-1}) LPS 4 h, and 0.1× (10 ng ml^{-1}) LPS 4 h. **k–n**, PCA of stimulated dendritic cells. **k**, First two principal components (or PCs) from a PCA performed on the LPS stimulation time course. From top to bottom: unstimulated/LPS 0 h, LPS 1 h, LPS 2 h, LPS 4 h, LPS 6 h. **l–n**, PCAs (left) and the distributions of scores (right) for each of the first three PCs for samples collected after stimulation with LPS (**l**), PAM (**m**), or PIC (**n**), for 1 (yellow), 2 (blue), 4 (grey) and 6 (red) hours. A single PCA was performed for all cells in all three time courses.

Extended Data Figure 2 | Effects of shallow read depth on expression estimates. a, b, A million reads per cell are sufficient to estimate expression levels. **a,** Scatter plot for a single cell (from Shalek *et al.*[16]) showing the relation between expression estimates calculated using 30 M reads (*x* axis) or a sub-sample of 1 M reads (*y* axis). **b,** Scatter plots for six different dendritic cells stimulated for 4 h with LPS. Each plot shows the relation between expression estimates calculated using all reads (*x* axis; number of reads marked on axis label) or a sub-sample of 1 M reads (*y* axis). In all cases, $R > 0.99$. Note that although, in principle, no gene should be estimated as present only in a subsample but not the full data set, this does occur for a very small number of genes (for example, four genes in cell 3), representing a nuanced technical error in RNA-seq estimation. Consider two expressed genes, *A* and *B*, from distinct loci, but with a short stretch of sequence identity. At low sequencing depth, if reads only map to the shared region, estimation tools, such as RSEM[36]

(used here), can guess erroneously which gene is expressed, such that additional sequencing depth can 'flip' the assignment of an uncertain read from gene *A* to gene *B*. These cases are extremely rare, and have a negligible effect. **c–e,** A million reads per cell are sufficient to estimate μ, σ², and α. Scatter plots showing the relation between α (**c**), μ (**d**), and σ² (**e**) values estimated using 10 M reads per cell (on average; *x* axis) or a sub-sample of 1 M reads per cell (*y* axis) from RNA-seq libraries prepared from individual bone-marrow-derived dendritic cells stimulated for 4 h with LPS. **c,** For almost all genes, 1 M reads are sufficient to estimate α. For a very small fraction (<0.1%) of weakly expressed genes, estimates of α are improved with increased sequencing depth. For μ (**d**) and σ² (**e**), estimates are plotted for all genes (left), only genes detected in more than 10 cells (middle), or only those genes detected in more than 30 cells (right).
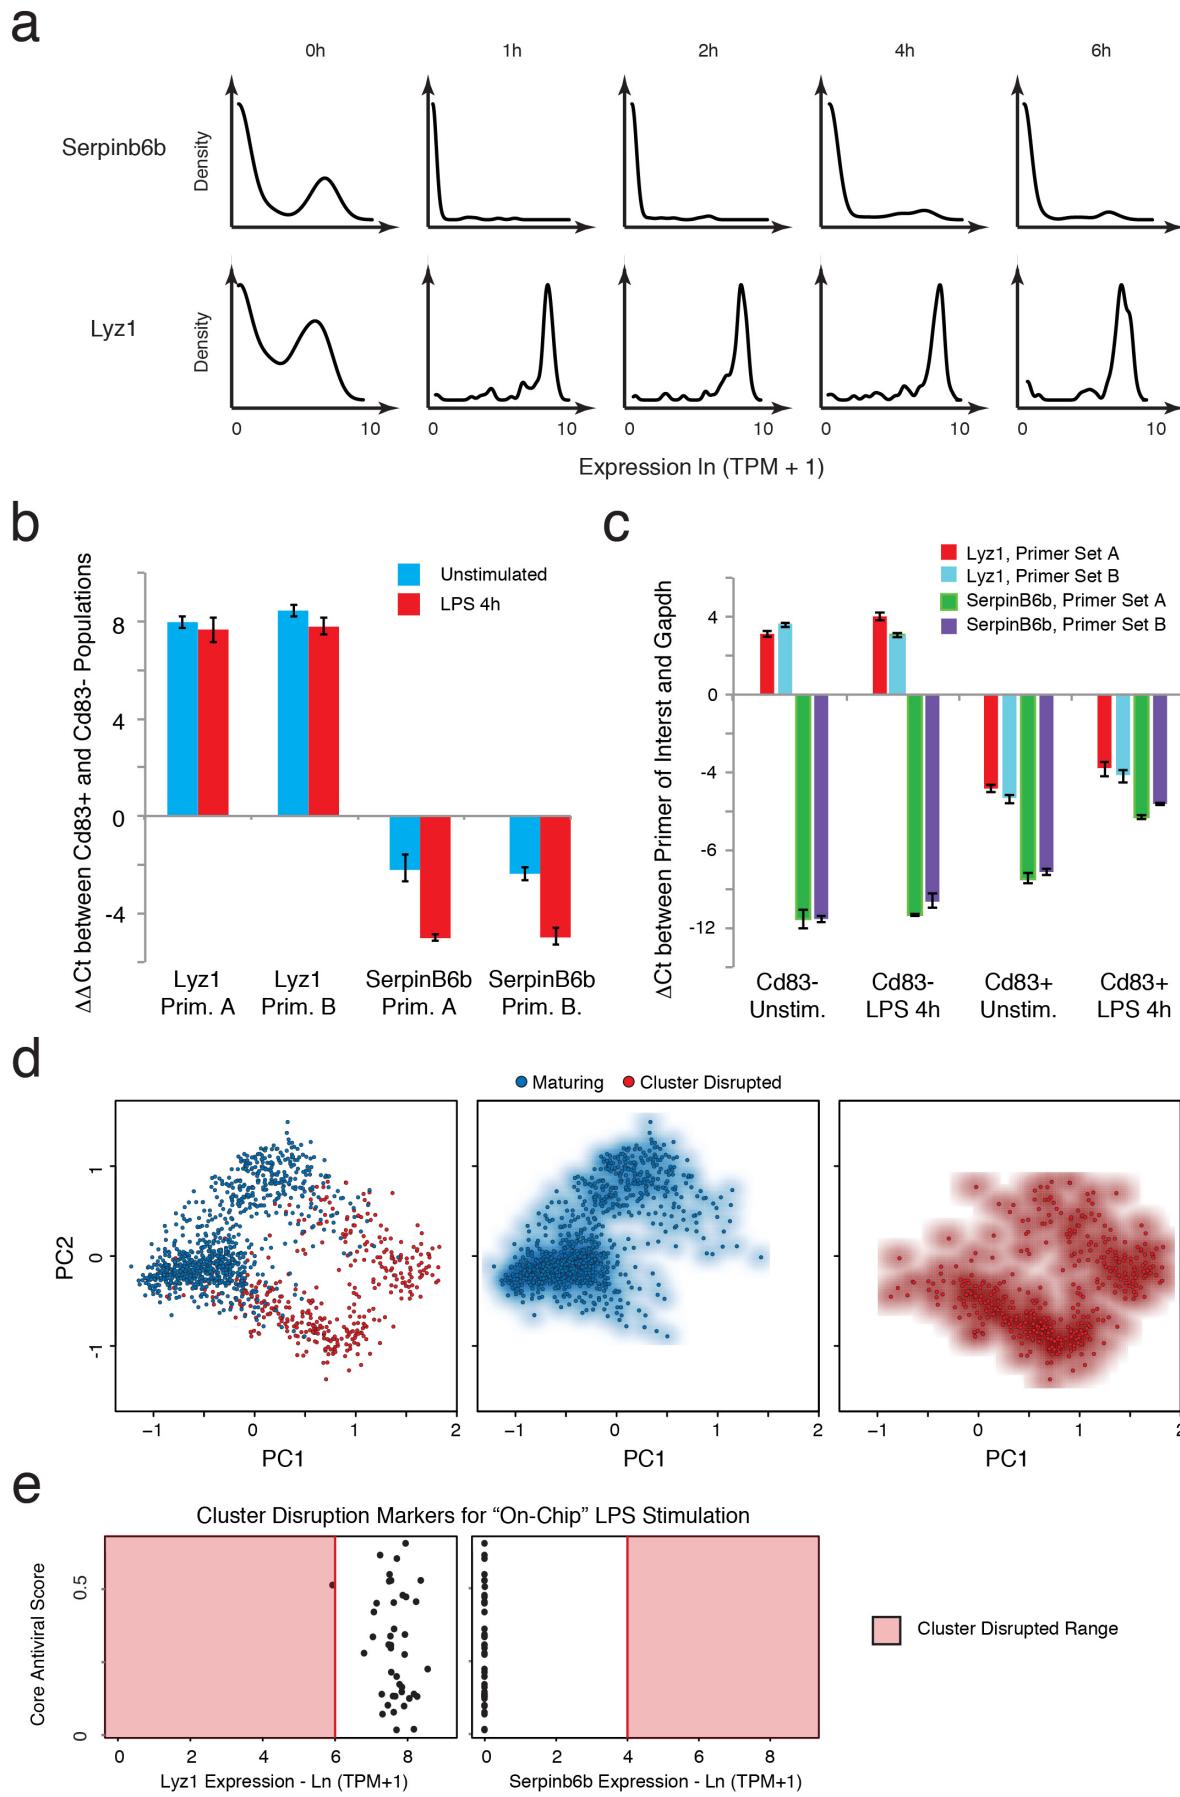
Extended Data Figure 3 | Technical and biological reproducibility.
a–d, Scatter plots showing the relationship between the average single-cell expression estimates in either of two technical replicates (LPS 2 h (**a**), LPS 4 h (**b**)) or two biological replicates (unstimulated/LPS 0 h (**c**), LPS 4 h (**d**)) for all genes (top), immune response genes (middle), or non-immune response genes (bottom). **e**, **f**, QQ plots (top) and MA plots (bottom) showing the similarity in expression for the two LPS 2 h technical replicates (**e**) or the two LPS 4 h technical replicates (**f**). Plots are provided across all genes (left), non-immune response genes (middle), or immune response genes (right). **g**, **h**, QQ plots (top) and MA plots (bottom) showing the similarity in expression estimates for all cells (including cluster-disrupted cells) in the two
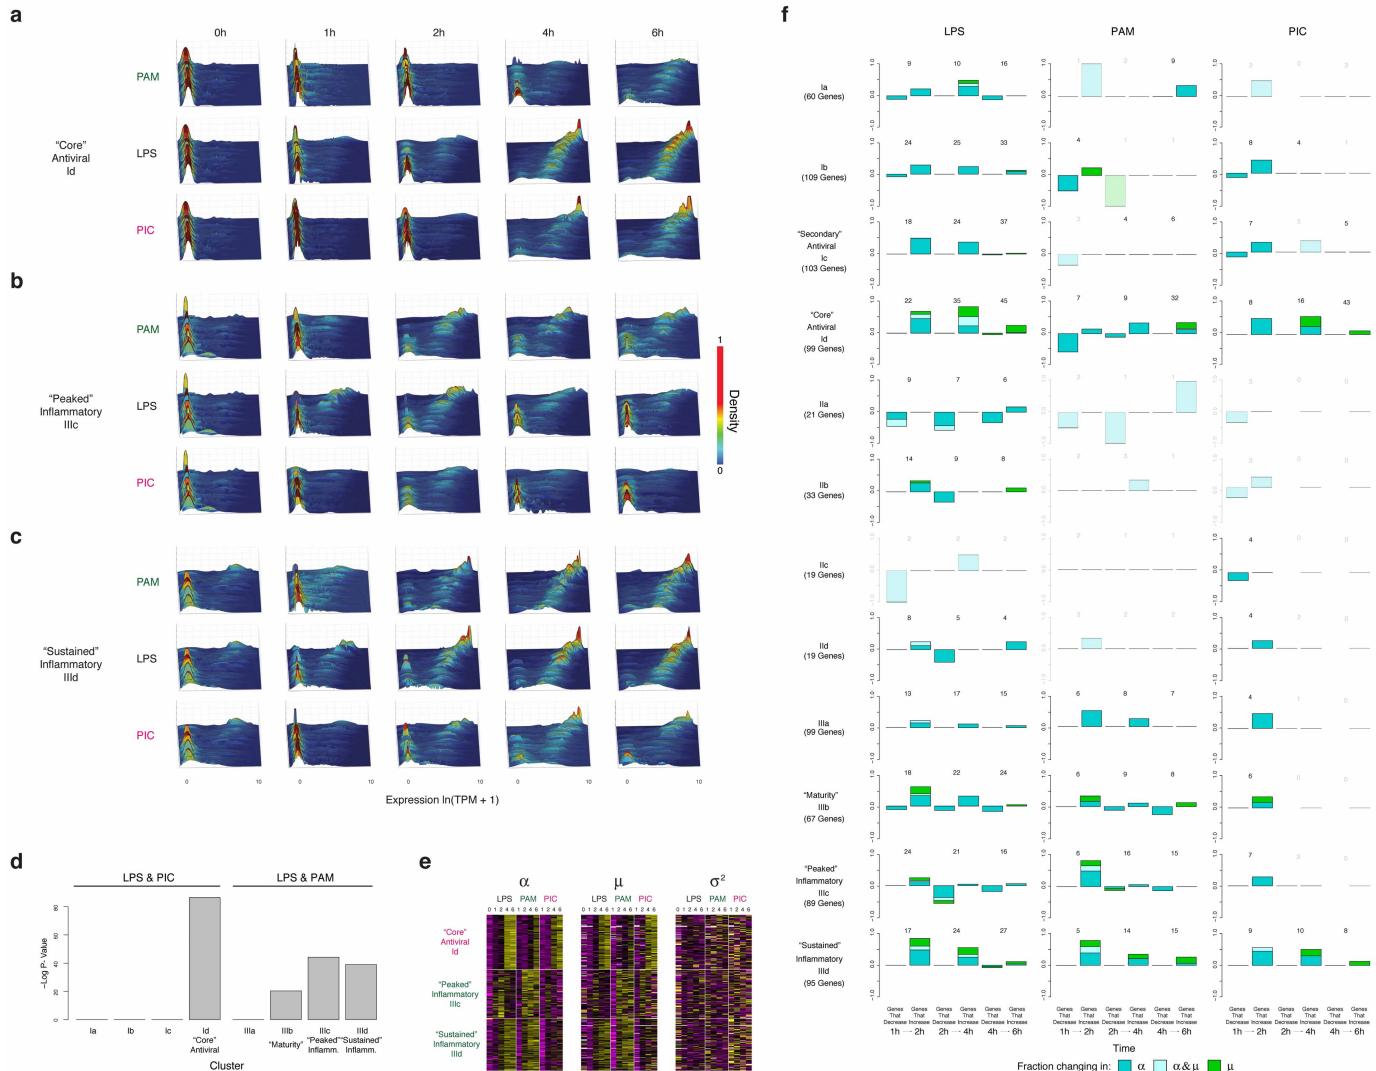
LPS 0 h/unstimulated biological replicates (**g**) or the two LPS 4 h biological replicates (**h**) across either all genes (left), non-immune response genes (middle), or immune response genes (right). Note, that slight variations in the fraction of cluster-disrupted cells and activation state of one of the two 0 h samples results in mild deviations between immune response gene estimates in those biological replicates. **i**, **j**, PCA for the two LPS 4 h technical replicates. **i**, The first two principal components (PC1 and PC2, x and y axis, respectively) of a PCA from the two LPS 4 h stimulation technical replicates (blue: replicate 1; red: replicate 2). **j**, The distributions of scores for cells from each of the two technical replicates on each of the first five PCs (left to right: PC1 to PC5).

a

Serpinb6b

Lyz1

0h 1h 2h 4h 6h

Density

Expression ln (TPM + 1)

b

ΔΔCt between Cd83+ and Cd83- Populations

Unstimulated
LPS 4h

Lyz1
Prim. A

Lyz1
Prim. B

SerpinB6b
Prim. A

SerpinB6b
Prim. B.

c

ΔCt between Primer of Interst and Gapdh

Lyz1, Primer Set A
Lyz1, Primer Set B
SerpinB6b, Primer Set A
SerpinB6b, Primer Set B

Cd83-
Unstim.

Cd83-
LPS 4h

Cd83+
Unstim.

Cd83+
LPS 4h

d

Maturing   Cluster Disrupted

PC2

PC1   PC1   PC1

e

Core Antiviral Score

Cluster Disruption Markers for "On-Chip" LPS Stimulation

Lyz1 Expression - Ln (TPM+1)

Serpinb6b Expression - Ln (TPM+1)

Cluster Disrupted Range

Extended Data Figure 4 | Cluster disruption. a, Single-cell expression distributions for *SerpinB6b* (a positive marker of cluster disruption) and *Lyz1* (a negative marker of cluster disruption) at each time point (marked on top) after stimulation with LPS (all cells included, see Supplementary Information). Distributions are scaled to have the same maximum height. **b**, Difference in mRNA expression as measured by qRT–PCR (with a *Gapdh* control) between *Lyz1* or *SerpinB6b* in cells pre-sorted before stimulation on the presence or absence of CD83 expression (CD83$^+$ and CD83$^-$, respectively), a known cell surface marker of cluster-disrupted cells (see Supplementary Information). Pre-sorted cells were then either unstimulated (blue) or stimulated (red) with LPS for 4 h. **c**, Expression of cluster-disruption markers does not change with stimulation. qRT–PCR showing the difference between *Gapdh* (control) and *Lyz1* or *SerpinB6b* expression in cells pre-sorted on Cd83 either in the presence or absence of simulation with LPS. **d**, PCA showing the separation between maturing' (blue dots) and cluster-disrupted (red dots) cells. **e**, Expression of cluster disruption markers for cells stimulated with LPS on-chip. For each cell (black dot) stimulated with LPS on-chip, shown are the expression levels (*x* axis) of *SerpinB6b* (cluster disruption cell marker, left) and *Lyz1* (normal maturing cell marker, right) versus its antiviral score (*y* axis). With one exception, the cells are clearly maturing and not cluster-disrupted. Red shading: range of expression in cluster-disrupted cells.
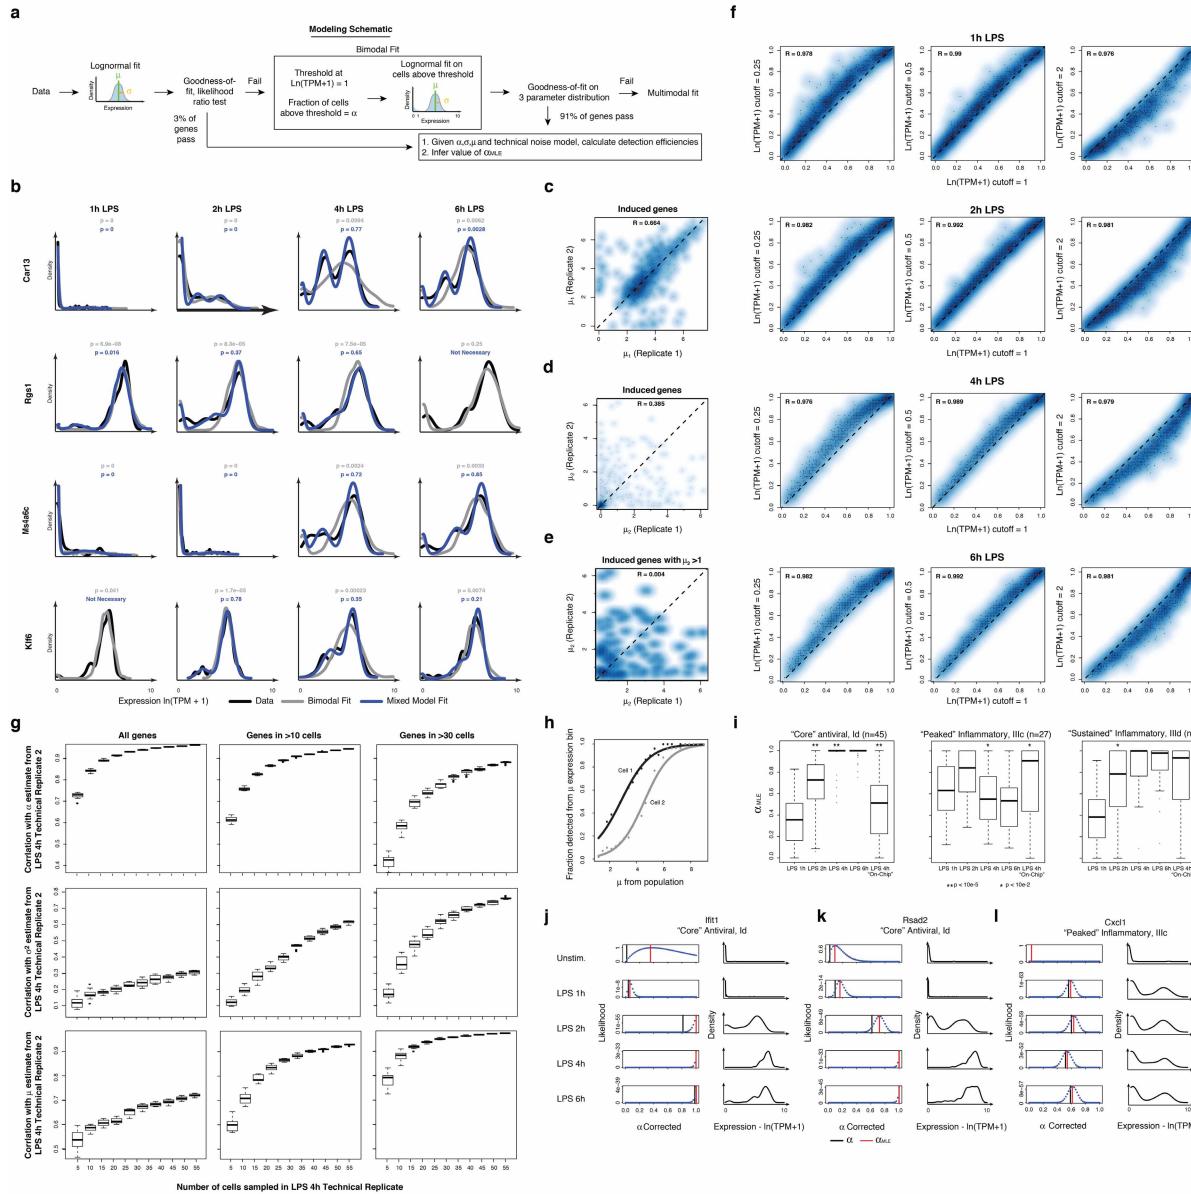
Extended Data Figure 5 | Time-dependent behaviours of single cells from different modules and stimuli. a–c, For each of the three key modules: core antiviral, I_d (**a**), peaked inflammatory, III_c (**b**), and sustained inflammatory, III_d (**c**) shown are wave plots of all of its constituent genes in bone-marrow-derived dendritic cells stimulated with PAM (top), LPS (middle), or PIC (bottom) for 0, 1, 2, 4 and 6 h (left to right). x axis: expression level, ln(TPM + 1); y axis: genes; z axis: single-cell expression density (proportion of cells expressing at that level). Genes are ordered from lowest to highest average expression value at the 4 h LPS time point. **d,** Contributions of each module to measured variation. Significance of the contribution of modules I_a–I_d and III_a–III_d from Fig. 1 to the variation measured throughout the stimulation time course. Shown is the P value (Mann–Whitney test) of the tested association between each gene module and the first three PCs, calculated using a statistical resampling method (see Supplementary Information). Only the core antiviral, maturity, and peaked/sustained inflammatory clusters show statistically significant enrichments with the three PCs. **e,** Gene modules show coherent shifts in single-cell expression. Shown are heat maps of scaled

α (left), μ (middle), and σ^2 (right) values (colour bar, bottom) in each time course (LPS, PAM, PIC) for the genes in each of the three key modules (rows, modules marked on left). Heat maps are row-normalized across all three stimuli, with separate scalings for each of the three parameters, to highlight temporal dynamics. Genes are clustered as in Fig. 1. **f,** Dynamic changes in variation during stimulation for each module. For each module (rows) and stimulus (columns), shown are bar plots of the fraction of genes (y axis) with a significant change only in α (by a likelihood ratio test, $P < 0.01$, blue), only in μ (Wilcoxon test, $P < 0.01$, green), or in both (each test independently, light blue), at each transition (x axis), in different conditions (marked on top), separated by whether they increase or decrease during that transition. In each module and condition, the proportion is calculated out of the genes in the module that are significantly bimodal (by a likelihood ratio test) in at least one time point during the LPS response and are expressed in at least 10 cells in both conditions. Their number is marked on top of each bar; conditions with 3 or fewer genes changing are semi-transparent.
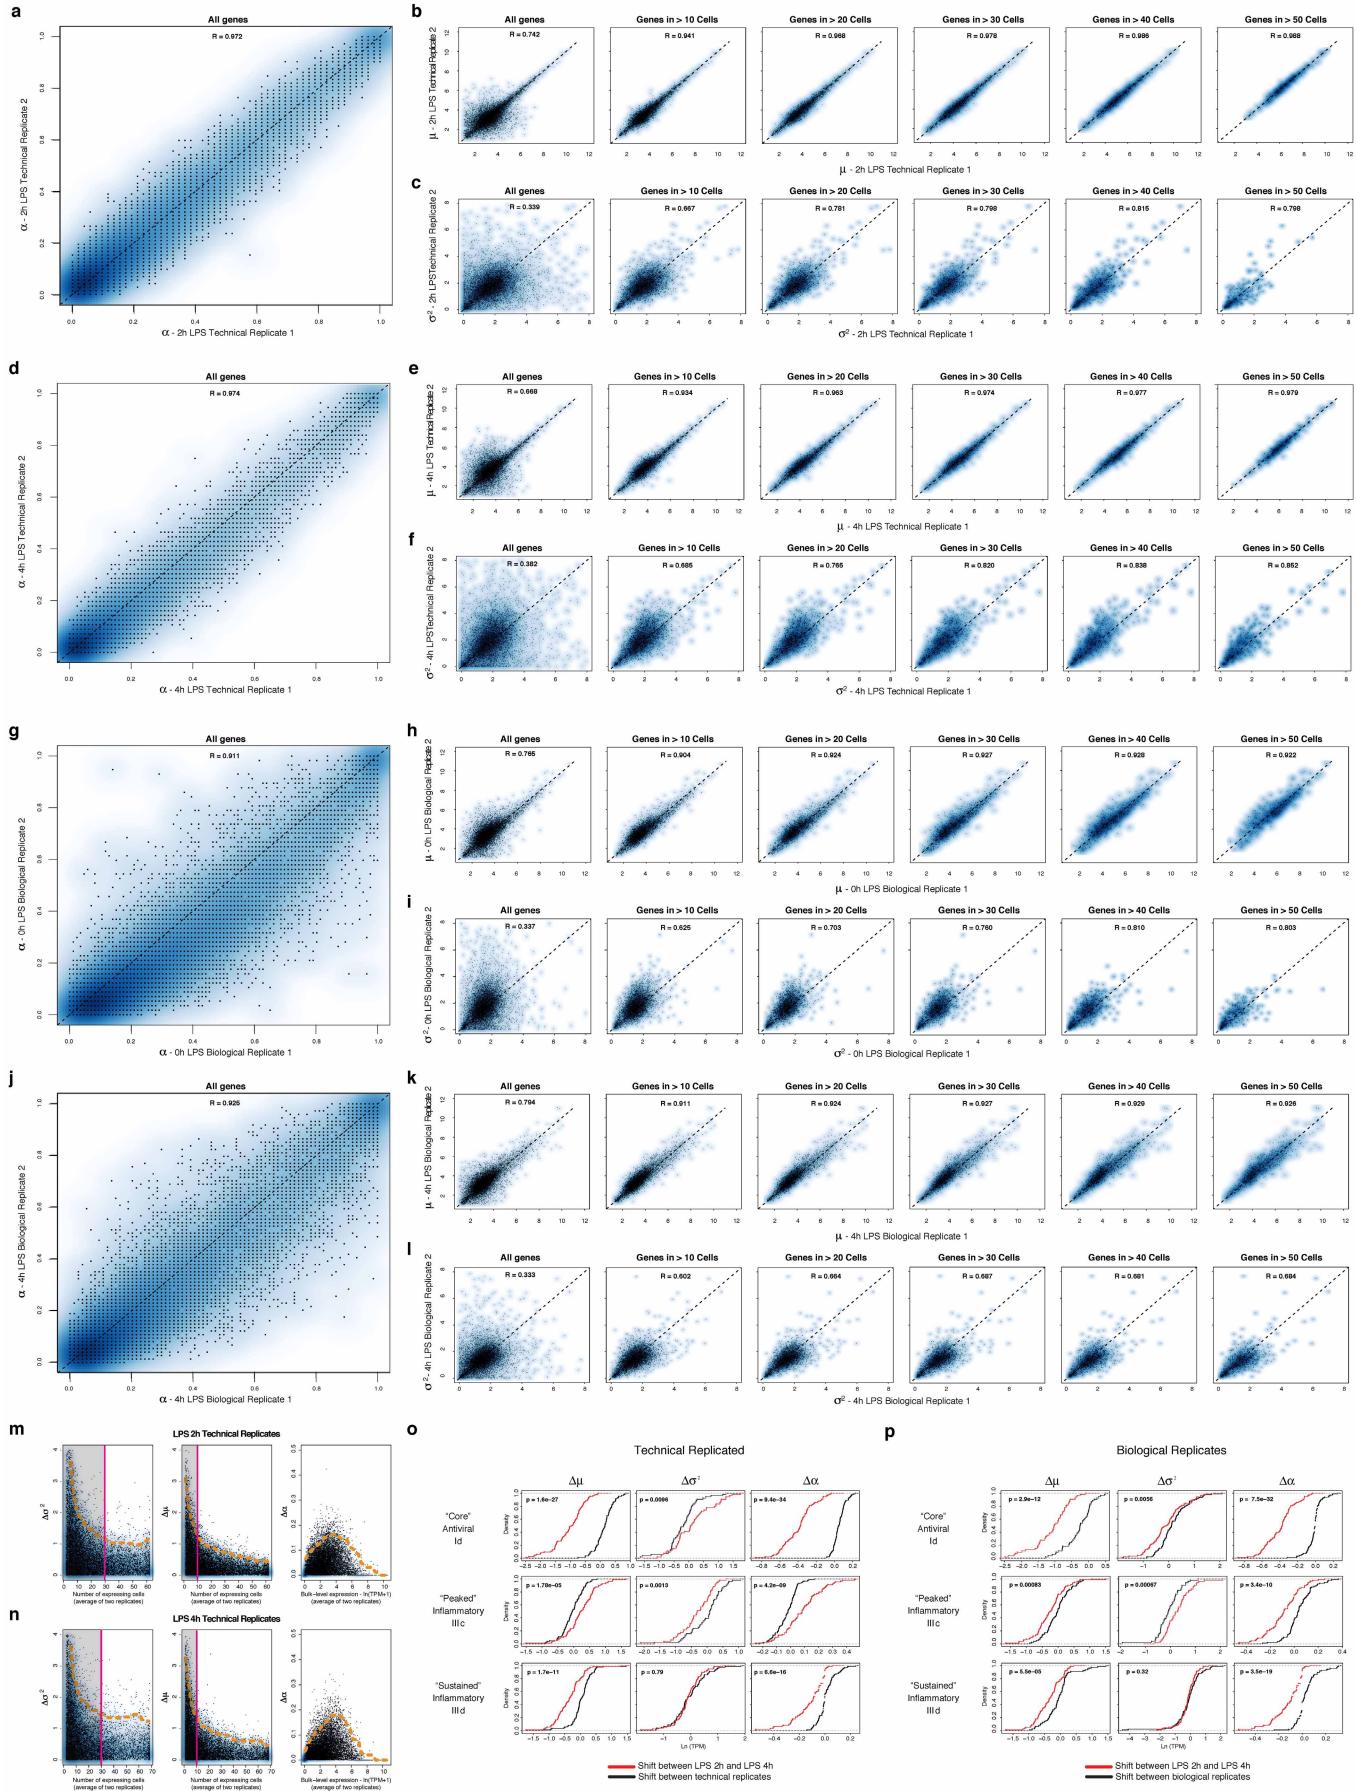
Extended Data Figure 6 | Comparison of single-cell RNA-seq to RNA-FISH. a–f, Single-cell mRNA expression distributions by RNA-FISH and single-cell RNA-seq. **a,** Representative images of genes analysed by RNA-FISH at 1 h and 4 h after LPS stimulation. **b–f,** mRNA expression distributions for the housekeeping gene *B2m* (**b**), the peaked inflammatory gene *Cxcl1* (**c**), the core antiviral gene *Ifit1* (**d**), the sustained inflammatory gene *Il6* (**e**), and the peaked inflammatory gene *Tnf* (**f**) measured using either single-cell RNA-seq (top, black curves) or RNA-FISH (black histograms; no smoothing) in either unstimulated cells (LPS 0 h) or cells stimulated with LPS for 1, 2, 4 or 6 h. **g–j,** Determining the detection limit of single-cell RNA-seq by comparison to RNA-FISH. For each of 25 genes, we compared single-cell RNA-seq data (*y* axis, this study) to RNA-FISH data (*x* axis, from Shalek *et al.*) based on either

their μ values (**g**, **h**) or α values (**i**, **j**), in which for RNA-FISH, expressing cells were defined for μ or α calculation at different thresholds (from left to right: at least 1, 4, 5, 10 or 20 copies per cell). **g, i,** Data from all 25 genes. **h, j,** Data after excluding probes from 5 cluster-disrupted gene markers (blue; *Ccl22*, *Ccr7*, *Irf8*, *SerpinB6b*, *Stat4*), which are less comparable since there are more cluster-disrupted cells in RNA-FISH experiments, and 2 low quality probes (grey; *Pkm2*, Fus) that showed very low expression in RNA-FISH, but had high expression in both single cell and population-level RNA-seq experiments. SPE (square-root of percent explained, top) represents the square-root of the total variance in the RNA-seq parameter explained by the RNA-FISH parameter, under the $y = x$ model compared to a constant fit (that is, $y = \bar{y}$).
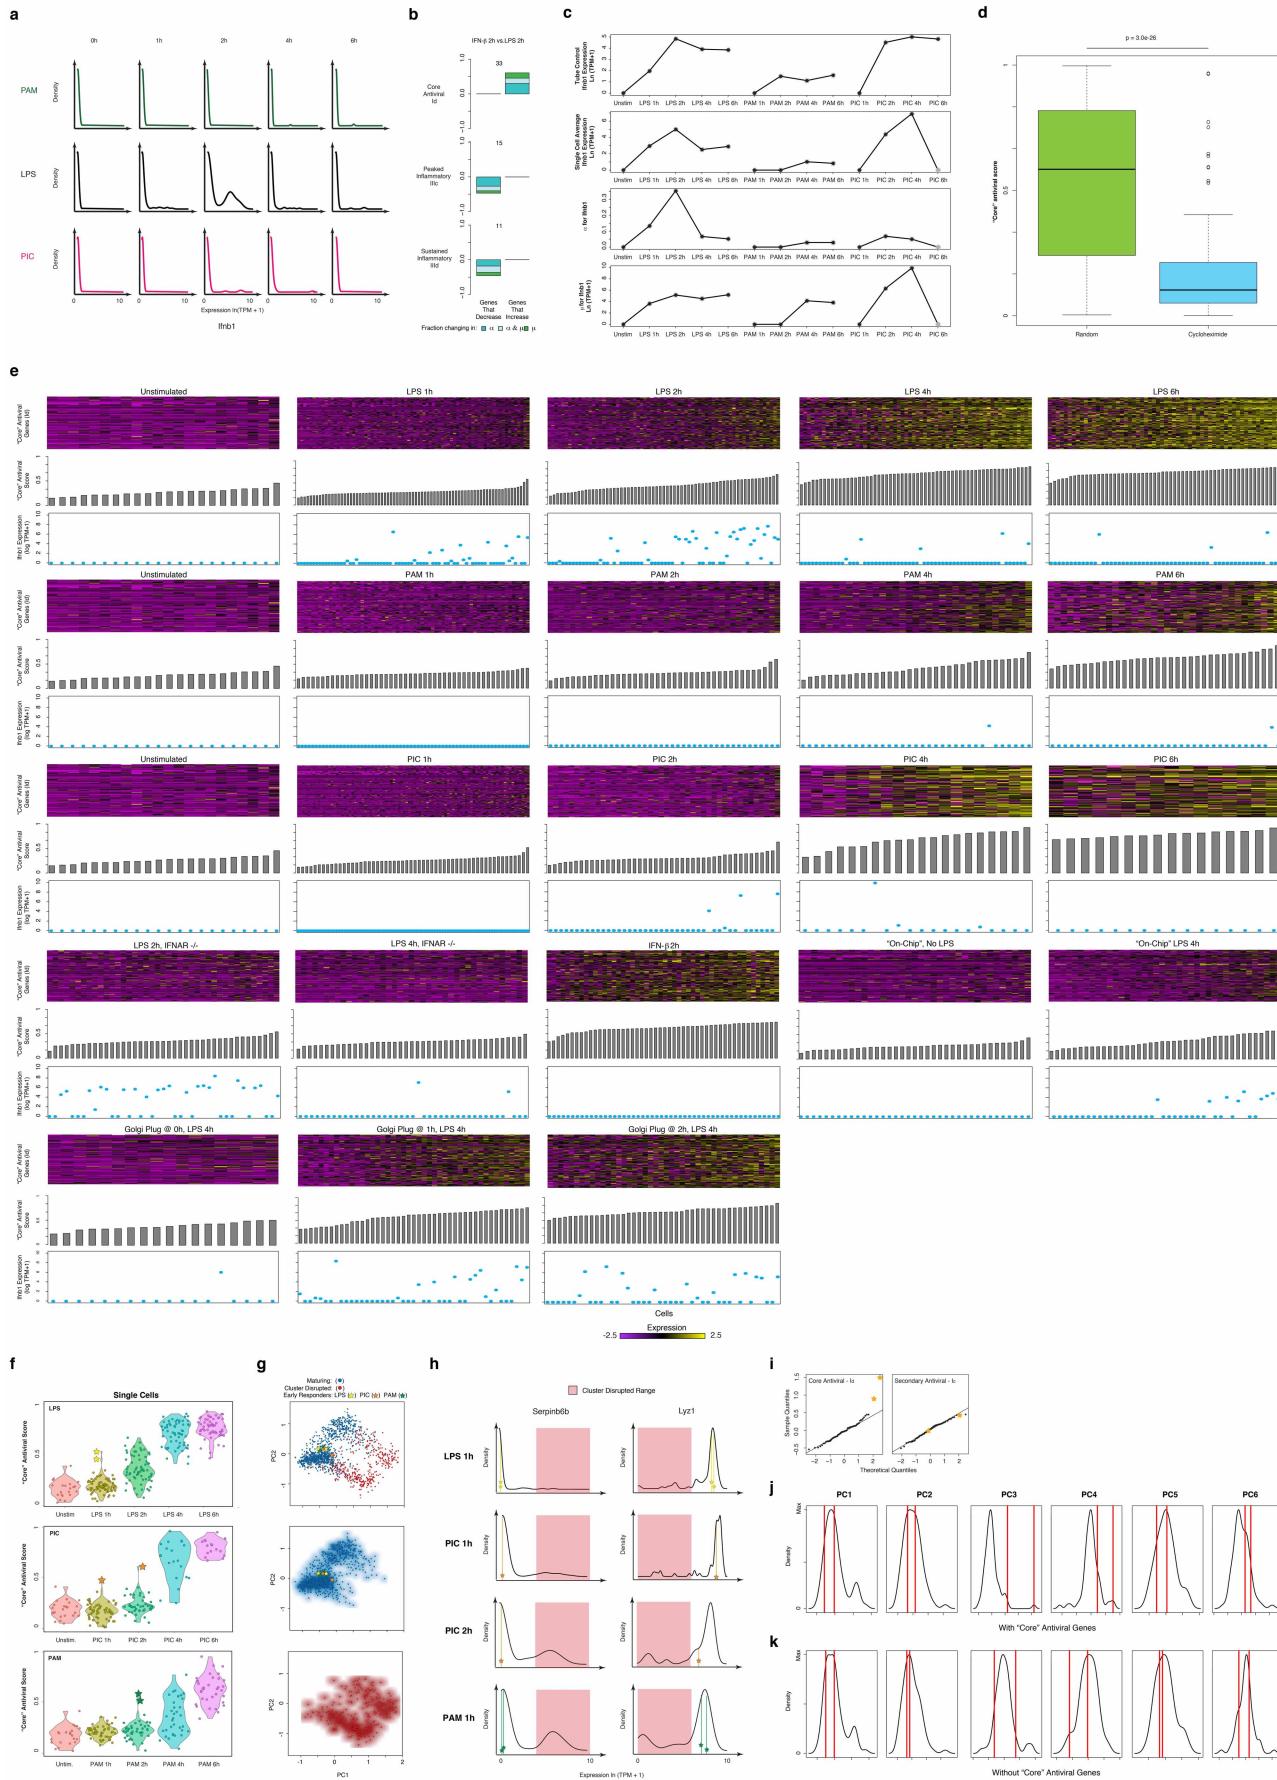
Extended Data Figure 7 | Fitting gene expression distributions. a, Flow chart of model fitting. Shown are the key steps in fitting the 3-parameter model. **b,** Examples of cases where fitting a multimodal distribution is required. Single-cell expression distributions for (top to bottom) *Car13*, *Rgs1*, Ms4a6c and *Klf6* (at left to right) 1, 2, 4, and 6 h (marked on top) after stimulation with LPS. Distributions are scaled to have the same maximum height. Data: black lines; Bimodal fits: grey lines; Multimodal fits: blue lines. *P* values (colour-coded) calculated using a goodness-of-fit test (a low *P* value rejects the fit; see Supplementary Information). **c–e,** Reproducibility of gene-specific fitting of the undetected mode, when fitting a mix of two normal distributions to all data points, including those with $\ln(\text{TPM} + 1) < 1$. **c, d,** Scatter plot showing the correlation between μ_1 and μ_2 estimates for the two LPS 4 h technical replicates (Supplementary Information), where μ_1 and μ_2 are the two component means (in decreasing order of magnitude) of the two mixed normal distributions. Estimates for μ_2 correlate poorly between technical replicates, particularly when focusing on genes for which μ_2 is greater than 1 (**e**), suggesting that the current data set does not support the use of this additional fit parameter. **f,** Robustness of α estimates to small deviations in the threshold. Scatter plots showing the correlation between α estimates determined when using a cut-off of $\ln(\text{TPM} + 1) = 1$ (*x* axis) versus when using a cut-off of $\ln(\text{TPM} + 1) = 0.25$ (*y* axis, left); 0.5 (*y* axis, middle) or 2 (*y* axis, right) for the LPS time course (top to bottom: 1 h, 2 h, 4 h and 6 h). **g,** Saturation curves for estimates of μ, σ^2, and α. Box plots depicting the Pearson correlation coefficient between α (top), μ (middle), or σ^2 (bottom) in two LPS 4 h technical replicates, as a function of

the number of cells randomly drawn from each replicate (full details in Supplementary Information). Plots are shown for all genes (left), as well as those detected in more than 10 (middle) or 30 cells, (right) in both replicates (full data sets). **h, i,** Correcting for the relationship between mean expression and average detection. **h,** The probability of detecting a transcript (*y* axis) in a cell as a function of μ (*x* axis). Black, grey curves are two illustrative cells from the LPS 4 h time point. **i,** Differences in α_{MLE}, a stringently-corrected MLE estimate of α (Supplementary Information), across the LPS time course. Shown are the box plots of α_{MLE} values (*y* axis) for bimodally expressed genes (determined by a likelihood ratio test, Supplementary Information) at each time point (1, 2, 4, and 6 h) following LPS stimulation (*x* axis), as well as for the on-chip 4 h LPS stimulation, for each of the core antiviral (left), peaked inflammatory (middle) and sustained inflammatory (right) modules. Stars represent intervals where there is a significant difference in a parameter between two consecutive time points, as determined by a Wilcoxon rank sum test (single star: $P < 10^{-2}$; double star: $P < 10^{-5}$). **j–l,** Estimating an upper bound on α using a likelihood test. For each of three transcripts (*Ifit1* (**j**); *Rsad2* (**k**); and *Cxcl1* (**l**)), shown are their expression distributions (red, left) and the matching likelihood function for a stringent upper estimate of α (blue dots, right), when considering a null model where expression is distributed in a lognormal fashion and any deviations are due to technical detection limits (Supplementary Information). Red vertical line: α_{MLE}; black vertical line: nominal α. Vertical green bars signify the nominal estimation of α, representing the fraction of cells with detected expression of a transcript.
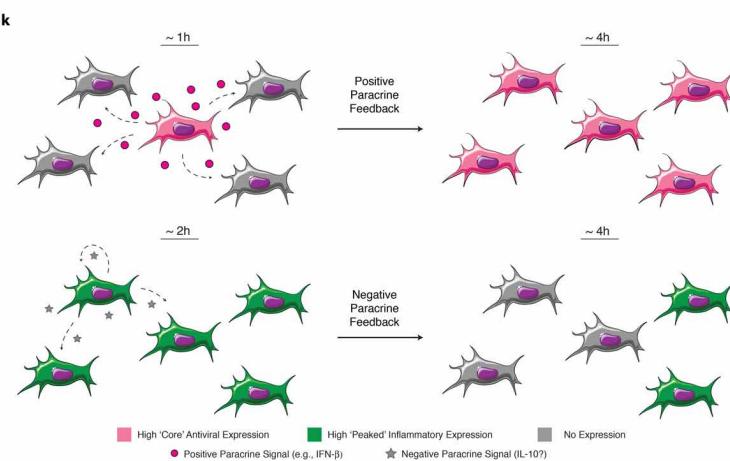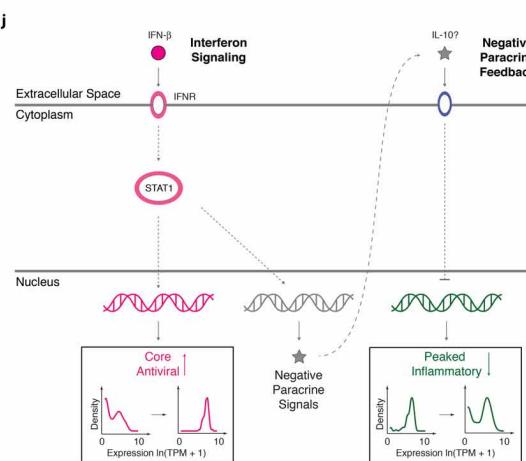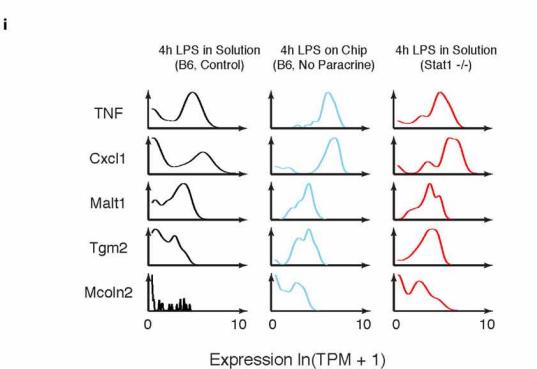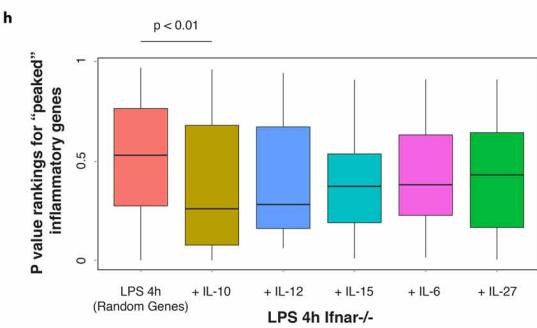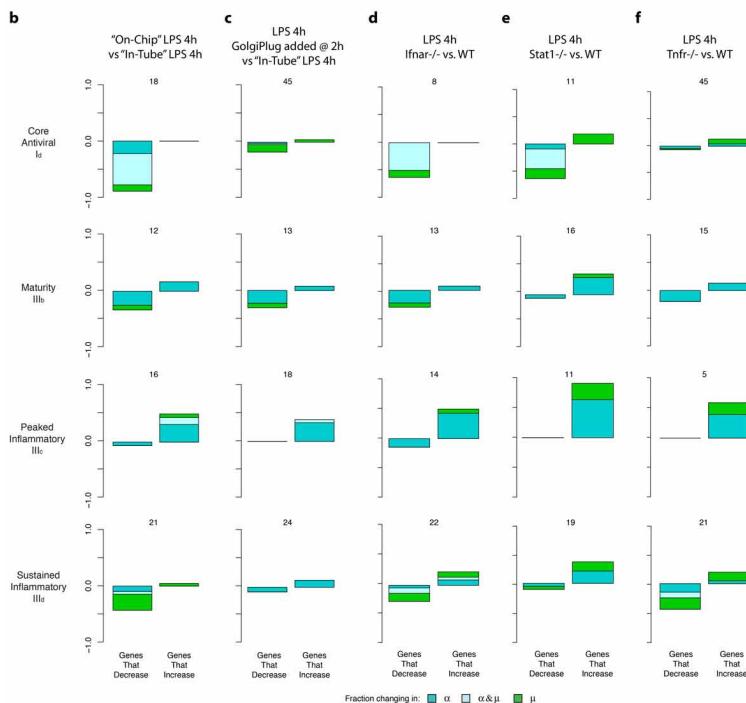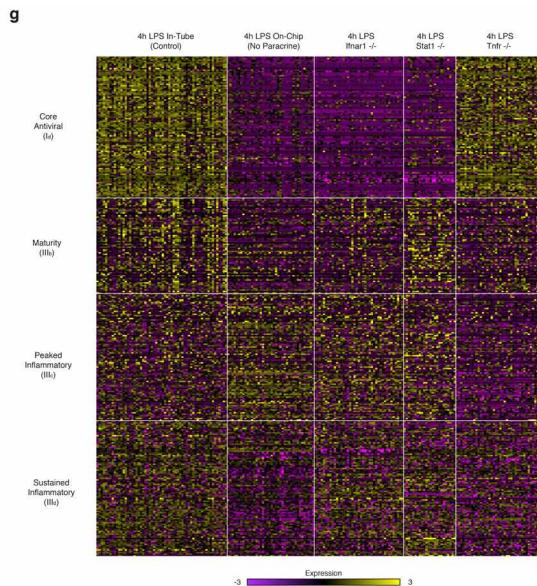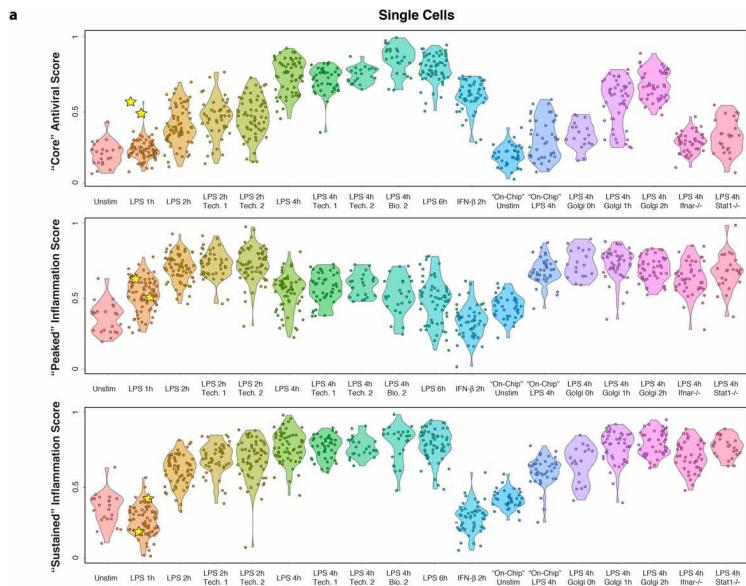
Extended Data Figure 8 | Reproducibility of estimated μ, σ² and α parameters. **a–f**, Reproducibility of estimated μ, σ² and α parameters between technical replicates. Scatter plots showing the relation between the estimated α (**a**), μ (**b**), and σ² (**c**) values for the two unstimulated/LPS 0 h technical replicates. For μ (**b**) and σ² (**c**), estimates are plotted for all genes (farthest on the left), as well as (left to right) for genes only detected in more than 10, 20, 30, 40 or 50 cells, respectively. **d–f**, show the same plots for the two LPS 4 h technical replicates. **g**, **h**, Reproducibility of estimated μ, σ² and α parameters between biological replicates. Scatter plots showing the relation between the α (**g**), μ (**h**), and σ² (**i**) values estimates for the two LPS 2 h biological replicates. For μ (**h**) and σ² (**i**), estimates are plotted for all genes (farthest on the left), as well as (left to right) for genes only detected in more than 10, 20, 30, 40 or 50 cells, respectively. **j–l**, show the same plots for the two LPS 4 h biological replicates. **m**, **n**, Relationship between per-gene errors for μ, σ² and α and the number of cells in which the gene's expression is detected, or its bulk expression level. Scatter plots displaying the differences in the σ² (left), μ (middle) and

α (right) estimates for each gene between technical replicates for LPS 2 h (**m**) or LPS 4 h (**n**) (y axis) as a function of either the number of cells in which the transcript is detected (x axis, for μ and σ²), or the transcript's bulk expression level (TPM, x axis, for α). Notably, σ² (left) estimates saturate (denoted by a magenta line and shaded box) after ~30 detectable events, whereas μ estimates saturate after ~10. Dashed orange line: 95% confidence interval. **o**, **p**, Changes in μ, σ² and α between time points are significant as compared to null models from both technical and biological replicates. Shown are the cumulative distribution functions (CDF) for shifts in μ (left), σ² (middle), and α (right) between 2 h and 4 h (red CDF) for the core antiviral (top), peaked inflammatory (middle), and sustained inflammatory (bottom) modules compared to a null model (black CDF) derived from either technical (**o**) or biological (**p**) replicates (Supplementary Information). In the vast majority of cases, the changes between time points are significant, as assessed by a Kolmogorov–Smirnof (KS) test (P value of test in the upper left corner of each plot).

Extended Data Figure 9 | *Ifnb1* expression, production, and precocious cells. a, b, *Ifnb1* mRNA expression and the effect of IFN-β on variation. **a**, Single-cell expression distributions for the *Ifnb1* transcript at each time point (top) after stimulation with PAM (top, green), LPS (middle, black), or PIC (bottom, magenta). Distributions are produced with the density function in *R* with default parameters, and scaled to have the same maximum density. **b**, For each of three modules defined in Fig. 1 (core antiviral, top; peaked inflammatory, middle; sustained inflammatory, bottom), shown are bar plots of the fraction of genes (*y* axis) with a significant change only in α (by a likelihood ratio test, $P < 0.01$, blue), only in μ (Wilcoxon test, $P < 0.01$, green), or in both (each test independently, light blue) between the 2 h LPS stimulation and the 2 h IFN-β stimulation separated by whether they increase or decrease during that transition. In each module and condition, the proportion is calculated out of the genes in the module that are significantly bimodal (by a likelihood ratio test) in at least time point during the LPS response and are expressed in at least 10 cells in both conditions. Their number is marked on top of each bar. **c, d**, *Ifnb1* mRNA expression patterns and effect of cycloheximide. **c**, From top to bottom, population average *Ifnb1* mRNA expression (top), single-cell average *Ifnb1* mRNA expression (second to top), α (second to bottom) and μ (bottom) estimates for *Ifnb1* for each stimulation condition in Fig. 1. Grey star at 6 h for PIC denotes uncertainty due to the small number of cells captured at that time point. **d**, Shown are box plots of the core antiviral scores (population level, see Supplementary Information) after a 4 h LPS stimulation either where cycloheximide was added at the time of stimulation (right, blue), or during a standard 4 h LPS control (left, green). Core antiviral expression is dramatically decreased by the addition of cycloheximide, suggesting that newly produced protein is needed to initiate the antiviral response. **e**, Relationship between core antiviral gene expression and *Ifnb1* mRNA expression during the LPS, PAM and PIC stimulation time courses and in follow-up experiments. Shown are the expression of core antiviral genes (heat maps: rows, gene; columns, cells) for cells stimulated for 0, 1, 2, 4 or 6 h (left to right) with LPS (top), PAM (middle), or PIC (bottom). Beneath each heat map, grey bars depict the core antiviral score (middle panel, see

Supplementary Information) and blue dots show *Ifnb1* mRNA expression for each cell in every heat map. **f–k**, Identifying the precocious cells. **f**, Core antiviral scores for cells stimulated with LPS, PIC, or PAM. Shown are violin plots of the core antiviral module scores (Supplementary Information, *y* axis) for each cell from time course experiments (from left: 0, 1, 2, 4 and 6 h) of dendritic cells stimulated with LPS (top), PIC (middle) or PAM (bottom). Two precocious cells (yellow stars, top panel) have unusually high antiviral scores at 1 h LPS (yellow stars, top); similarly precocious cells can be seen in PIC at 1 h and 2 h (orange stars, middle) or in PAM at 2 h (turquoise stars, bottom). **g**, Precocious cells in all three responses are typical maturing cells. PCA showing the separation between maturing (blue dots) and cluster-disrupted (red dots) cells (top), as well as only maturing (middle) or only cluster-disrupted (bottom) cells (all as also shown in Extended Data Fig. 4d). The precocious cells from each of the responses are marked as stars (colours as in (**f**)), and all fall well within the maturing cells. **h**, Precocious cells in all three responses express *Lyz1* and do not express *SerpinB6b*. Shown are mRNA expression distributions for *SerpinB6b* (cluster disruption cell marker, left) and *Lyz1* (normal maturing cell marker, right) in LPS 1 h, PIC 1 h and 2 h, and PAM 1 h (top to bottom). The typical range for expression in cluster-disrupted cells is shaded in red. The precocious cells of the responses are marked as stars (colours as in (**f**)), and all fall well within the maturing cells. **i**, Normal quantile plots of the expression of genes from the core (cluster I_d, left) and secondary (cluster I_c, right) antiviral clusters at 1 h LPS. The two precocious cells (yellow stars) express unusually high levels of core antiviral genes (left) but not of secondary genes (right). **j, k**, The precocious cells are only distinguished by the expression of ~100 core antiviral genes. Shown are the distributions of scores for each of the first six PCs (right) for samples collected after stimulation with LPS for 1 h with (**j**) or without (**k**) the inclusion of core antiviral genes. Precocious cells (vertical red bars), normally distinguished by the third and fourth principle components (**j**), become indistinguishable from all other cells if the ~100 core antiviral genes are excluded (**k**) before performing the PCA.

Extended Data Figure 10 | Characterizing the precocious cells. a, Core antiviral, peaked inflammatory, and sustained inflammatory module scores during the LPS time course and follow-up experiments. Shown are violin plots of the scores (y axis) for the core antiviral (Supplementary Information, top), peaked inflammatory (Supplementary Information, middle), and sustained inflammatory (Supplementary Information, bottom) modules for cells in each of the experiments (from left to right): LPS 0 h, LPS 1 h, LPS 2 h, LPS 2 h technical replicate 1, LPS 2 h technical replicate 2, LPS 4 h, LPS 4 h technical replicate 1, LPS 4 h technical replicate 2, LPS 4 h biological replicate, LPS 6 h, IFN-β 2 h, on-chip unstimulated, on-chip LPS 4 h, LPS 4 h with GolgiPlug at 0 h, LPS 4 h with GolgiPlug at 1 h, LPS 4 h with GolgiPlug at 2 h, LPS 4 h with $Ifnar^{-/-}$ dendritic cells, and LPS 4 h with $Stat1^{-/-}$ dendritic cells. Yellow stars: the two precocious cells at 1 h LPS. **b,** Changes in expression and variation during stimulation in the on-chip 4 h LPS stimulation. For genes in the (from top to bottom) core antiviral, maturity, peaked inflammatory and sustained inflammatory modules, shown are bar plots of the fraction of genes (y axis) with a significant change only in α (by a likelihood ratio test, $P < 0.01$, blue), only in μ (Wilcoxon test, $P < 0.01$, green), or in both (each test independently, light blue) between the 4 h on-chip LPS stimulation and the 4 h in-tube LPS stimulation separated by whether they increase or decrease during that transition. In each module and condition, the proportion is calculated out of the genes in the module that are significantly bimodal (by a likelihood ratio test) in at least one time point during the LPS response and are expressed in at least 10 cells in both conditions. Their number is marked on top of each bar. **c–f,** Bar plots, as in **b,** for a 4 h wild-type LPS stimulation (in-tube) and either a 4 h in-tube LPS stimulation where GolgiPlug was added 2 h after LPS (**c**), a 4 h LPS stimulation of $Ifnar^{-/-}$ dendritic cells (**d**), a 4 h LPS stimulation of $Stat1^{-/-}$ dendritic cells (**e**), or a 4 h LPS Stimulation of $Tnfr^{-/-}$ dendritic cells (**f**).

g, Genetic perturbations alter expression and variation in different inflammatory and antiviral modules. Shown is the expression of the genes (rows) in, from top to bottom: the core antiviral (I_d), maturity (III_b), peaked inflammatory (III_c), and sustained inflammatory (III_d) modules in single cells (columns) in, from left to right: the in-tube, on-chip, $Ifnar1^{-/-}$, $Stat1^{-/-}$, and $Tnfr^{-/-}$ conditions. Yellow/purple colour scale: scaled expression values (z-scores). **h,** Scores of the peaked inflammatory module for $Ifnar^{-/-}$ dendritic cells with recombinant cytokines. Shown are box plots of the peaked inflammatory module scores (Supplementary Information, y axis) for three population-level replicates of a 4 h LPS stimulation of $Ifnar^{-/-}$ dendritic cells to which a recombinant cytokine (x axis) has been added at 2 h after stimulation. Notably, adding IL-10 significantly reduces the peaked inflammatory module. **i,** $Stat1$ knockout affects expression and variation of peaked inflammatory genes. Shown are expression distributions for five peaked inflammatory genes after 4 h of LPS stimulation in each of three conditions: in-tube stimulation of wild-type dendritic cells (control; left), on-chip stimulation of wild-type cells (no cell-to-cell signalling; middle), and a stimulation of dendritic cells from $Stat1^{-/-}$ mice (performed in-tube; right). **j, k,** Population-level paracrine signalling enhances and coordinates the core antiviral response while dampening and desynchronizing the peaked inflammatory ones. **j,** Gene network model showing how positive IFN-β signalling induces the antiviral response and reduces its heterogeneity, while simultaneously activating negative paracrine feedback, possibly including IL-10, which dampens the peaked inflammatory cluster and increases its heterogeneity. **k,** Cell population model showing how positive and negative paracrine signalling alter antiviral (magenta) and inflammatory (green) gene expression variability across cells. Grey denotes no expression.